FortiOS™ Handbook - Networking

VERSION 5.6.3

FORTIOS 5.6 VERSION

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

FORTICAST

http://forticast.fortinet.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FORTINET PRIVACY POLICY

https://www.fortinet.com/corporate/about-us/privacy.html

FEEDBACK

Email: techdocs@fortinet.com

4/26/2018

FortiOS™ Handbook - Networking

01-563-456038-20180426

TABLE OF CONTENTS

Change log 9
About this guide 10
What's new in Networking 11
FortiOS version 5.6.3 11
VXLAN loopback binding support 11
New option to configure DHCP renew time 11
Fixed inability to delete multicast interfaces in the GUI 11
New CLI commands for showing and adding IPv6 addresses 11
Static route GUI updates 11
DHCP relay agent option configuration 12
FortiOS version 5.6.2 12
FortiOS version 5.6.1 12
Recursive DNS server option 12
Routing Monitor support for policy routes 12
Control how routing changes affect active sessions 12
FortiOS version 5.6 12
New command to view transceiver information 13
BGP local-AS support 13
Interface setting removed from SNMP community configuration page 13
Remove RPF checks from the state evaluation process 13
BGP enhancements 13
FQDN support for static routes 13
Priority for blackhole routes 13
DDNS refresh interval 14
GUI support for configuring IPv6 blackhole routes 14
Support for SSL VPN and WAN link load balancing 14
DDNS support for No-IP 14
IPv6 Router Advertisement options for DNS configuration 14
SD-WAN replaces WAN LLB in FortiGate GUI 14
Interfaces 15
Administrative access 15
Aggregate interfaces 16
Sending GARP on aggregate MAC changes 17
DHCP addressing mode on an interface 17
DHCP servers and relays 19
Configuring DHCP servers 19
Configuring the DHCP relay agent option 21
Configuring DHCP with IPv6 22
Service 23

Configuring the lease time 23
Configuring the DHCP renew time 23
DHCP options 24
Excluding addresses in DHCP 24
Viewing information about DHCP server connections 25
Breaking an address lease 25
Interface MTU packet size 25
Interface settings 26
Interface configuration and settings 28
Loopback interfaces 32
VXLAN loopback binding 32
One-armed sniffer 32
Ports preassigned as sniffer ports 34
Physical ports 35
Configuring the FortiGate 100D ports 35
Displaying information about the status of transceivers 35
Split port support 36
PPPoE addressing mode on an interface 36
Probing interfaces 38
Enhanced TWAMP Light functionality with server/controller functionality 39
Redundant interfaces 39
Dual Internet connections 39
Redundant interfaces 40
Load sharing 42
Secondary IP addresses to an interface 42
Software switch 43
Soft switch example 44
Virtual switch 44
Support for 802.1x fallback and 802.1x dynamic VLANs 45
Zones 47
Virtual domains 48
Wireless 49
VLANs 50
VLANs in NAT mode 51
Adding VLAN subinterfaces 52
Configuring security policies and routing 54
VLANs in transparent mode 55
VLANs and transparent mode 55
General configuration steps 55
VLAN switching and routing 57
VLAN layer-2 switching 57
VLAN layer-3 routing 58

Layer-2 and ARP traffic 58
STP forwarding 59
STP support for FortiGate models with hardware switches 59
ARP traffic 59
Multiple VDOMs solution 60
Vlanforward solution 60
Forward-domain solution 61
Asymmetric routing 61
NetBIOS 62
Too many VLAN interfaces 63
Troubleshooting VLAN issues 63
Botnet and command-and-control protection 63
DNS 65
DNS settings 65
Additional DNS CLI configuration 65
DDNS 66
Configuring FortiGate to refresh DDNS IP addresses 66
TLS support for DDNS updates 67
DDNS update override for DHCP 67
FortiDDNS registration to a public IP address 67
DNS servers 67
Configuring a recursive DNS 69
Advanced static routing 72
Routing concepts 72
Routing in VDOMs 72
Default route 72
Adding a static route 73
Enabling or disabling individual static routes 73
Configuring FQDNs as a destination address in static routes 73
Routing table 74
Configuring the maximum number of IP route cache entries 78
Building the routing table 81
Static routing security 82
Multipath routing and determining the best route 84
Route priority 85
Use of firewall addresses for static route destinations 86
Removing RPF checks from the state evaluation process 86
Troubleshooting static routing 87
Static routing tips 88
Policy routing 89
Adding a policy route 90
Enabling or disabling individual policy routes 92

Moving a policy route 93
Use of firewall addresses for policy route destinations 94
Transparent mode static routing 94
Static routing example 95
Network layout and assumptions 95
General configuration steps 96
Get your ISP information such as DNS, gateway, etc 96
Configure FortiGate unit 96
Configure administrator PC and dentists' PCs 101
Testing network configuration 102
Dynamic routing overview 104
What is dynamic routing? 104
Comparing static and dynamic routing 104
Dynamic routing protocols 105
Minimum configuration for dynamic routing 106
Comparison of dynamic routing protocols 107
Features of dynamic routing protocols 107
When to adopt dynamic routing 110
Choosing a routing protocol 112
Dynamic routing terminology 113
Controlling how routing changes affect active sessions 118
IPv6 in dynamic routing 119
RIP 120
RIP background and concepts 120
Background 120
RIP terminology and parts 121
How RIP works 125
Troubleshooting RIP 131
Routing loops 131
Holddowns and triggers for updates 134
Split horizon and poison reverse updates 135
Debugging IPv6 on RIPng 135
Simple RIP example 136
Network layout and assumptions 136
General configuration steps 138
Configuring FortiGate system information 138
Configuring FortiGate unit RIP router information 147
Configuring other networking devices 150
Testing network configuration 151
RIPng: RIP and IPv6 151
Network layout and assumptions 152
Configuring the FortiGate units system information 153

Configuring RIPng on FortiGate units 156
Configuring other network devices 157
Testing the configuration 157
OSPF 158
OSPF background and concepts 158
Background 158
The parts and terminology of OSPF 158
How OSPF works 165
Troubleshooting OSPF 170
Clearing OSPF routes from the routing table 170
Checking the state of OSPF neighbors 170
Passive interface problems 170
Timer problems 171
BFD 171
Authentication issues 171
DR and BDR election issues 171
Basic OSPF example 172
Network layout and assumptions 172
Configuring the FortiGate units 174
Configuring OSPF on the FortiGate units 177
Configuring other networking devices 184
Testing network configuration 184
Advanced inter-area OSPF example 184
Network layout and assumptions 184
Configuring the FortiGate units 186
Configuring OSPF on the FortiGate units 190
Configuring other networking devices 195
Testing network configuration 195
Controlling redundant links by cost 195
Adjusting the route costs 197
Verifying route redundancy 198
BGP 199
BGP background and concepts 199
Background 199
Parts and terminology of BGP 199
How BGP works 209
Troubleshooting BGP 213
Clearing routing table entries 213
Route flap 213
Dual-homed BGP example 217
Network layout and assumptions 218
Configuring the FortiGate unit 220

Configuring other networking devices 228
Testing this configuration 229
Redistributing and blocking routes in BGP 230
Network layout and assumptions 230
Configuring the FortiGate unit 231
Testing network configuration 236
IS-IS 238
IS-IS background and concepts 238
Background 238
How IS-IS works 238
Parts and terminology of IS-IS 240
Troubleshooting IS-IS 244
Routing loops 244
Split horizon and poison reverse updates 247
Simple IS-IS example 247
Network layout and assumptions 248
Expectations 249
CLI configuration 249
Verification 251
Troubleshooting 254
Multicast forwarding 255
Sparse mode 255
Dense mode 257
PIM support 257
Multicast forwarding and FortiGate units 258
Multicast forwarding and RIPv2 258
Configuring FortiGate multicast forwarding 259
Adding multicast security policies 259
Enabling multicast forwarding 260
Displaying IPv6 multicast router information 262
Multicast routing examples 262
Example FortiGate PIM-SM configuration using a static RP 262
FortiGate PIM-SM debugging examples 271
Example multicast DNAT configuration 276
Example PIM configuration that uses BSR to find the RP 278
Troubleshooting 287
Netflow support 287
sFlow support 287
Configuration 288

Change log

Date	Change description
January 24, 2018	• Updated with information about FortiOS 5.6.3 features • Moved sFlow support information from the *System Administration Handbook* to this handbook
December 20, 2017	FortiOS 5.6.2 release

About this guide

This guide explains how to configure your network. It contains the following sections:

Chapter title	Description
Interfaces	• Explains the concepts of options for setting up interfaces and groupings of sub-networks that can scale to a company's growing requirements
DNS	• How to set DNS requirements for your network • How to set FortiGate as a local DNS server
Advanced static routing	• Explains universal and static routing concepts, equal cost multipath (ECMP) and load balancing, policy routing, and routing in transparent mode
Dynamic routing overview	• Provides an overview of dynamic routing, compares static and dynamic routing, and helps you decide which dynamic routing protocol is best for you
RIP	• Describes a distance-vector routing protocol intended for small, relatively homogeneous networks
OSPF	• Provides background on the specific protocol explaining terms used and how the protocol works, as well as providing some troubleshooting information and examples on configuring the protocols in different situations
BGP	• Classless inter-domain routing • Aggregate routes • BGP is the only routing protocol to use TCP for a transport protocol
IS-IS	• Describes the link state protocol, is well-suited to smaller networks and with near universal support on routing hardware. The section also provides troubleshooting information and configuration examples
Multicast forwarding	• Concepts and use of multicasting with the FortiGate
Troubleshooting	• Describes features, such as packet capture, that are useful for troubleshooting purposes

What's new in Networking

This section contains a list of new features and enhancements for Networking.

FortiOS version 5.6.3

FortiOS version 5.6.3 includes the following new features and enhancements for Networking.

VXLAN loopback binding support

A Virtual Extensible LAN (VXLAN) unicast device can bind to a loopback interface as its underlying interface.

For more information, see VXLAN loopback binding on page 32.

New option to configure DHCP renew time

You can now set a minimum DHCP renew time.

For more information, see Configuring the DHCP renew time on page 23.

Fixed inability to delete multicast interfaces in the GUI

An issue where you could not delete interfaces in the GUI under **Network > Multicast > Multicast Routing** has been fixed.

New CLI commands for showing and adding IPv6 addresses

In previous releases, you could only view client IPv6 DHCP addresses in the CLI by using `diagnose ipv6 address list`. This release introduces the ability to also view these addresses by entering `get` in the interface (under `config system interface`).

Similarly, in previous releases, you could only view static IPv6 addresses under `diagnose ipv6 address list` and by entering `get` under the interface. This release introduces the ability to also view them by entering `fnsysctl/sysctl ifconfig interface_id`.

Static route GUI updates

The GUI pages for creating or editing static routes (under **Network > Static Routes**) have been updated.

When you create or edit a static route, you must specify the gateway first and then select the device interface. The FortiGate GUI will try to populate the **Interface** field based on the value in the **Gateway** field. If you specify a gateway that is not in the same subnet as the selected interface, a warning message will appear.

Also, the **Device** field has been renamed to the **Interface** field.

DHCP relay agent option configuration

FortiGate now supports the ability to enable or disable the DHCP relay agent option.

For more information, see Configuring the DHCP relay agent option on page 21.

FortiOS version 5.6.2

There are no new features for Networking in FortiOS version 5.6.2.

FortiOS version 5.6.1

FortiOS version 5.6.1 includes the following new features and enhancements for Networking.

Recursive DNS server option

FortiGate now supports the following RFC 6106 IPv6 Router Advertisement options:

- Sending DNS search list option to downstream clients with Router Advertisements that use a static prefix
- Sending Recursive DNS server option to downstream clients with Router Advertisements that use a static prefix

For more information, see Configuring IPv6 Router Advertisement options for DNS configuration on page 69.

Routing Monitor support for policy routes

You can now monitor policy routes using the FortiGate GUI. The **Routing Monitor** page now includes a **Policy** option that lists the active policy routes on the FortiGate and provides information about them.

For more information, see Viewing the routing table on page 74.

Control how routing changes affect active sessions

You can now control how active sessions are affected when dynamic routing changes occur that affect the routes the active sessions are using. You can configure whether FortiGate maintains the original routing for the sessions that are using the affected routes, or applies the routing table changes to the active sessions.

For more information, see Controlling how routing changes affect active sessions on page 118.

FortiOS version 5.6

FortiOS version 5.6 includes the following new features and enhancements for Networking.

New command to view transceiver information

You can now use a command to display information about transceivers installed in FortiGate SFP/SFP+ interfaces. You can use this command on most FortiGate models that have SFP/SFP+ interfaces.

For more information, see Displaying information about the status of transceivers on page 35.

BGP local-AS support

FortiGate now supports BGP local-AS.

For more information, see BGP local-AS support on page 212.

Interface setting removed from SNMP community configuration page

The Interface column has been removed from the **New SNMP Community** configuration page (**System > SNMP > SNMP v1/v2c > Create New**) in the FortiGate GUI.

Remove RPF checks from the state evaluation process

You can now remove RPF (reverse path forwarding) state checks without needing to enable asymmetric routing, using the new `set src-check disable` CLI command.

For more information, see Removing RPF checks from the state evaluation process on page 86.

BGP enhancements

FortiGate now supports the following BGP enhancements:

- Option to stop BGP graceful restart process on timer only
- Option to bring down BGP neighbor upon link down
- Option to keep routes for a period after the BGP neighbor is down

For more information, see Configuring BGP graceful restart process on timer on page 211,Configuring option to bring down BGP neighbor when the link is down on page 212, and Configuring option to keep routes for a period after the BGP neighbor is down on page 212.

FQDN support for static routes

You can now configure FQDN firewall addresses as destination addresses in a static route.

For more information, see Configuring FQDNs as a destination address in static routes on page 73.

Priority for blackhole routes

You can now add a priority to a blackhole route to change its position relative to kernel routes in the routing table.

For more information, see Adding a blackhole route with a priority on page 84.

DDNS refresh interval

You can now configure FortiGate to refresh DDNS IP addresses by periodically checking the configured DDNS server.

For more information, see Configuring FortiGate to refresh DDNS IP addresses on page 66.

GUI support for configuring IPv6 blackhole routes

You can now configure IPv6 blackhole routes in the FortiGate GUI.

For more information, see Configuring IPv6 blackhole routes on page 83.

Support for SSL VPN and WAN link load balancing

You can now set virtual WAN link interfaces as destination interfaces in firewall policies for WAN link load balancing, when SSL VPN is the source interface.

For more information, see SSL VPN and WAN link load balancing on page 42.

DDNS support for No-IP

FortiGate now supports No-IP as a DDNS server.

IPv6 Router Advertisement options for DNS configuration

FortiGate now supports the following RFC 6106 IPv6 Router Advertisement options:

- Obtaining DNS search list options from upstream DHCPv6 servers
- Sending the DNS search list through Router Advertisement
- Sending the DNS search list through the FortiGate DHCP server

For more information, see Configuring IPv6 Router Advertisement options for DNS configuration on page 69.

SD-WAN replaces WAN LLB in FortiGate GUI

The term SD-WAN has replaced the term WAN LLB throughout the FortiGate GUI.

Interfaces

Interfaces, both physical and virtual, allow traffic to flow between internal networks and the Internet, and between internal networks. FortiGate has a number of options for setting up interfaces and groupings of subnetworks that can scale to your organization's growing requirements.

Administrative access

To help prevent FortiGate interfaces, especially the public-facing ports, from being accessed by users who you do not want accessing them, you can configure protocols that an administrator must use to access FortiGate, including:

- HTTPS
- PING
- FortiManager Access (FMG-Access)
- CAPWAP
- SSH
- SNMP
- FTM
- RADIUS Accounting
- FortiTelemetry

As a best practice, you should configure administrative access when you are setting the IP address for the port.

The following example adds an IPv4 address 172.20.120.100 to the WAN1 interface, and administrative access to HTTPS and SSH.

Add an IP address to the WAN1 interface - web-based manager

1. Go to **Network > Interface**.
2. Select the WAN1 interface row and select **Edit**.
3. Select the **Addressing Mode** of **Manual**.
4. Enter the IP address for the port of 172.20.120.100/24.
5. For **Administrative Access**, select **HTTPS** and **SSH**.
6. Select **OK**.

Add an IP address to the WAN1 interface - CLI

```
config system interface
   edit wan1
      set ip 172.20.120.100/24
      set allowaccess https ssh
end
```

When you add or remove a protocol, you must type the entire list of protocols again. For example, if you have an access list of HTTPS and SSH and you want to add PING, you must use the following CLI command:

```
set allowaccess https ssh ping
```

If you use `set allowsaccess ping`, only ping is set and HTTPS and SSH are removed.

Aggregate interfaces

Link aggregation (IEEE 802.3ad) allows you to bind two or more physical interfaces together to form an aggregated link. This new link has the bandwidth of all the links combined. If a link in the group fails, traffic is automatically transferred to the remaining interfaces with the only noticeable effect being reduced bandwidth.

This is similar to redundant interfaces, with the major difference being that a redundant interface group uses only one link at a time, while an aggregate link group uses the total bandwidth of the functioning links in the group, up to eight (or more).

Some FortiGate models support the IEEE standard 802.3ad for link aggregation.

An interface can be an aggregate interface if it meets the following criteria:

- It is a physical interface, not a VLAN interface or subinterface.
- It is not already part of an aggregate or redundant interface.
- It is in the same VDOM as the aggregated interface. Aggregate ports cannot span multiple VDOMs.
- It does not have an IP address and is not configured for DHCP or PPPoE.
- It is not referenced in any security policy, VIP, IP pool, or multicast policy.
- It is not an HA heartbeat interface.
- It is not one of the backplane interfaces of the FortiGate 5000 series.

Some FortiGate models do not support aggregate interfaces. In this case, the aggregate option is not available in the FortiGate GUI or CLI. Also, you cannot create aggregate interfaces from interfaces in a switch port.

To see if a port is being used or has other dependencies, use the following CLI command:

```
diagnose sys checkused system.interface.name <interface_name>
```

When an interface is included in an aggregate interface, it is not listed in the **Network > Interface** page in the FortiGate GUI. Interfaces still appear in the CLI, but if you configure those interfaces, it will not take effect. You cannot configure the interface individually and it is not available to include in security policies, VIPs, IP pools, or routing.

To avoid unintentional network issues when you configure Link Aggregation Control Protocol (LACP), disconnect the interfaces that you want to add to the aggregate interface. After you finish configuring LACP, reconnect the interfaces.

The following example creates an aggregate interface on a FortiGate 3810A, using ports 4 to 6, with an internal IP address of 10.13.101.100, and administrative access to HTTPS and SSH.

Create an aggregate interface - web-based manager

1. Go to **Network > Interface** and select **Create New**, then **Interface**.
2. Enter the Name as `Aggregate`.
3. For the **Type**, select **802.3ad Aggregate**.
   If this option does not appear, your FortiGate unit does not support aggregate interfaces.
4. In the **Physical Interface Members**, click to add interfaces. Select port 4, 5, and 6.
5. Select the **Addressing Mode** of **Manual**.
6. Enter the IP address for the port of 10.13.101.100/24.
7. For **Administrative Access**, select HTTPS and SSH.
8. Select **OK**.

Create aggregate interface - CLI

```
config system interface
   edit Aggregate
      set type aggregate
      set member port4 port5 port6
      set vdom root
      set ip 172.20.120.100/24
      set allowaccess https ssh
   end
```

Sending GARP on aggregate MAC changes

FortiGate sends out GARP (Gratuitous Address Resolution Protocol) announcements if the MAC address of a link aggregated interface changes to a new IP pool address due to a link failure or change in ports. This is needed when you use networking devices, such as some switches that do not perform this function when they receive LACP (Link Aggregation Control Protocol) information about changes in the MAC information.

DHCP addressing mode on an interface

If you configure an interface to use DHCP, FortiGate automatically broadcasts a DHCP request from the interface. The interface is configured with the IP address, any DNS server addresses, and the default gateway address that the DHCP server provides.

DHCP IPv6 is similar to DHCP IPv4, except:

• No default gateway option is defined because a host learns the gateway using router advertisement messages.
• There are no WINS servers because it is obsolete.

For more information about DHCP IPv6, see RFC 3315.

You can configure DHCP for an interface in **Network > Interfaces** in the FortiGate GUI. Select the interface from the list, and select **DHCP** in the **Addressing mode**. The following table describes the DHCP status information when DHCP is configured for an interface.

Field	Description
Status	Displays DHCP status messages as the interface connects to the DHCP server and gets addressing information. Select **Status** to refresh the addressing mode status message. Status can be one of the following values: • **initializing**: no activity • **connecting**:interface attempts to connect to the DHCP server • **connected**:interface retrieves an IP address, netmask, and other settings from the DHCP server • **failed**:interface was unable to retrieve an IP address and other settings from the DHCP server
Obtained IP/Netmask	The IP address and netmask leased from the DHCP server. This is only displayed if the **Status** is **connected**.
Renew	Select this to renew the DHCP license for this interface. This is only displayed if the **Status** is **connected**.
Expiry Date	The time and date when the leased IP address and netmask is no longer valid for the interface. The IP address is returned to the pool to be allocated to the next user request for an IP address. This is only displayed if the **Status** is **connected**.
Default Gateway	The IP address of the gateway defined by the DHCP server. This is displayed only if the **Status** is **connected**, and if **Receive default gateway from server** is selected.
Distance	Enter the administrative distance for the default gateway retrieved from the DHCP server. The administrative distance is an integer from 1 to 255, and specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route.
Retrieve default gateway from server	Enable this to retrieve a default gateway IP address from the DHCP server. The default gateway is added to the static routing table.
Override internal DNS	Enable this to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page. When VDOMs are enabled, you can override the internal DNS only on the management VDOM.

DHCP servers and relays

A DHCP server provides an address, from a defined address range, to a client on the network that requests it.

An interface cannot provide both a server and a relay for connections of the same type (regular or IPsec). However, you can configure a regular DHCP server on an interface only if the interface is a physical interface with a static IP address. You can configure an IPsec DHCP server on an interface that has either a static or a dynamic IP address.

You can configure one or more DHCP servers on any FortiGate interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

If an interface is connected to multiple networks through routers, you can add a DHCP server for each network. The IP range of each DHCP server must match the network address range. The routers must be configured for DHCP relay.

You can configure a FortiGate interface as a DHCP relay. The interface forwards DHCP requests from DHCP clients to an external DHCP server and returns the responses to the DHCP clients. The DHCP server must have appropriate routing so that its response packets to the DHCP clients arrive at the unit.

DHCP server options are not available in transparent mode.

Configuring DHCP servers

To add a DHCP server, go to **Network > Interfaces**. Edit the interface, and select **DHCP** in the addressing mode.

Field	Description
Address Range	By default, the FortiGate unit assigns an address range based on the address of the interface for the complete scope of the address. For example, if the interface address is 172.20.120.230, the default range created is 172.20.120.231 to 172.20.120.254. Select the range and select **Edit** to adjust the range or select **Create New** to add a different range.
Netmask	Enter the netmask of the addresses that the DHCP server assigns.
Default Gateway	Select this to use either **Same as Interface IP** or select **Specify** and enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
DNS Server	Select this to use **Same as system DNS**, **Same as Interface IP** or select **Specify** and enter the IP address of the DNS server.

Field	Description
Mode	Select the type of DHCP server FortiGate will be. By default, it is a **Server**. Select **Relay** if needed. When **Relay** is selected, the above configuration is replaced by a field to enter the **DHCP Server IP** address.
DHCP Server IP	This appears only when **Mode** is **Relay**. Enter the IP address of the DHCP server where FortiGate obtains the requested IP address.
Type	Select this to use the DHCP in **Regular** or **IPsec** mode.
Additional DHCP Options	Use this to create new DHCP options.
MAC Address + Access Control	Select this to match an IP address from the DHCP server to a specific client or device using its MAC address. In a typical situation, an IP address is assigned ad hoc to a client, and that assignment times out after a specific time of inactivity from the client, known as the lease time. To ensure a client or device always has the same IP address (there is no lease time), use IP reservation.
Add from DHCP Client List	If the client is currently connected and using an IP address from the DHCP server, you can select this option to select the client from the list.

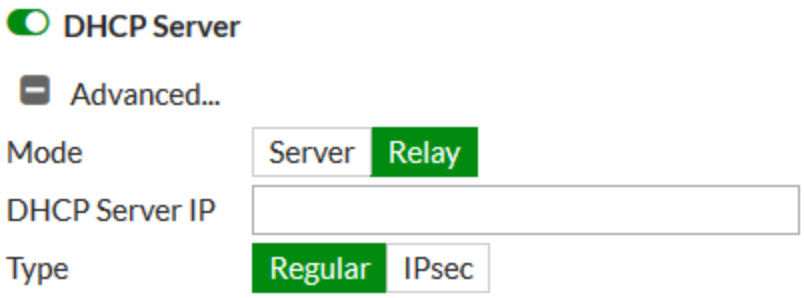

DHCP Server

Address Range

Create New	Edit	Delete
Starting IP	End IP	
192.168.10.100	192.168.10.254	

Netmask 255.255.255.0

Default Gateway Same as Interface IP | Specify

DNS Server Same as System DNS | Same as Interface IP | Specify

FortiClient On-Net Status

Advanced...

Mode Server | Relay

NTP Server Local | Same as System NTP | Specify 0.0.0.0

Time Zone Same as System | Specify

Next Bootstrap Server 0.0.0.0

Additional DHCP Options

Create New	Edit	Delete	
Seq #	Option Code	Value	Hexadecimal Value
	51 (Lease Time)	604800	

MAC Reservation + Access Control

Create New	Edit	Delete	Add from DHCP Client List
MAC Address	Action or IP	Description	
Unknown MAC Addresses	Assign IP		

Type Regular | IPsec

Configuring the DHCP relay agent option

You can configure the DHCP relay agent option (option 82 in RFC 3046). You can enable or disable whether the DHCP relay agent option is added. This option is disabled, by default.

To configure the DHCP relay agent option, use the following CLI commands:

```
config system interface
   edit <name>
      set dhcp-relay-agent-option [enable | disable]
   next
```

For more information about the DHCP relay option, see RFC 3046 (DHCP Relay Agent Information Option).

Configuring DHCP with IPv6

You can use DHCP with IPv6, using the CLI. To configure DHCP, ensure IPv6 is enabled by going to **System > Feature Visibility** and enable **IPv6** under **Basic Features**. Use the following CLI command:

```
config system dhcp6 server
```

For more information about the configuration options, see the FortiOS CLI Reference Guide.

DHCPv6 prefix delegation

Prefix delegation is supported for DHCP for IPv6 addressing. It is not practical to manually provision networks on a large scale in IPv6 networking. The DHCPv6 prefix delegation feature is used to assign a network address prefix, and automate the configuration and provisioning of the public routable addresses for the network.

You can enable the prefix delegation, using the following CLI commands:

```
config system interface
   edit "wan1"
      config ipv6
         set ip6-mode dhcp
         set ip6-allowaccess ping
         set dhcp6-prefix-delegation enable
         end
      end
```

DHCPv6 prefix hint

This feature is used to "hint" to upstream DCHPv6 servers a desired prefix length for their subnet to be assigned in response to its request.

There is a possibility of duplicate prefixes being sent by ISP when using a /64 bit subnet because the first 64 bits of the address are derived from the MAC address of the interface. This could cause an issue if the system administrator wishes to divide the host networks into 2 /64 bit subnets.

By receiving a /60 bit (for example) network address, the administrator can then divide the internal host works without the danger of creating duplicate subnets.

Also included in the new feature, are preferred times for the life and valid life of the DHCP lease.

DHCPv6 hint for the prefix length:

`set dhcp6-prefix-hint` <DHCPv6 prefix that will be used as a hint to the upstream DHCPv6 server>

DHCPv6 hint for the preferred life time:

`set dhcp6-prefix-hint-plt` <integer> 1 ~ 4294967295 seconds or "0" for unlimited lease time

DHCPv6 hint for the valid life time:

`set dhcp6-prefix-hint-vlt` <integer> 1 ~ 4294967295 seconds or "0" for unlimited lease time

Service

On low-end FortiGate units, a DHCP server is configured on the internal interface, by default, with the following values:

Field	Value
Address Range	192.168.1.110 to 192.168.1.210
Netmask	255.255.255.0
Default Gateway	192.168.1.99
Lease Time	7 days
DNS Server 1	192.168.1.99

These settings are appropriate for the default internal interface IP address of 192.168.1.99. If you change this address to a different network, you need to change the DHCP server settings to match.

Alternatively, after the FortiGate unit assigns an address, you can go to **Monitor > DHCP Monitor** and locate the specific user. Right-click and select **Create/Edit IP Reservation**.

Configuring the lease time

The lease time determines the length of time an IP address remains assigned to a client. Once the lease expires, the address is released for allocation to the next client that requests an IP address.

To configure the lease time, use the following CLI commands:

```
config system dhcp server
   edit <server_entry_number>
      set lease-time <seconds>
   next
end
```

The default lease time is seven days. To have an unlimited lease time, set the value to zero.

Configuring the DHCP renew time

You can set a minimum DHCP renew time. This option is available only when `mode` is set to `dhcp`.

To set the DHCP renew time, using the following CLI commands:

```
config system interface
   edit <name>
      set mode dhcp
      set dhcp-renew-time <seconds>
   next
end
```

The possible values for `dhcp-renew-time` are 300 to 605800 seconds (five minutes to seven days). To use the renew time that the server provides, set this entry to 0.

DHCP options

When you add a DHCP server, you can include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you may need to configure a FortiGate DHCP server that gives out a separate option, as well as an IP address, such as an environment that needs to support PXE boot with Windows images.

The option numbers and codes are specific to a particular application. The documentation for the application should provide the values you should use. Option codes are represented in option value and HEX value pairs. The option is a value between 1 and 255.

You can add up to three DHCP code/option pairs per DHCP server.

To configure option 252 with value http://192.168.1.1/wpad.dat (FortiGate CLI)

```
config system dhcp server
   edit <server_entry_number>
      set option1 252 687474703a2f2f3139322e3136382e312e312f777061642e646174
end
```

For more information about DHCP options, see RFC 2132 (DHCP Options and BOOTP Vendor Extensions).

FortiGate DHCP works with DDNS to allow FQDN connectivity to leased IP addresses

As clients are assigned IP addresses, they send back information that would be found in an A record to the FortiGate DHCP server, which can take this information and pass it back to a corporate DNS server so that even devices using leased IP address can be reached using FQDNs. You can configure the settings for this feature using the `ddns-update` CLI command and some other ddns related options.

DHCP server option fields

In place of specific fields, the DHCP server maintains a table for the potential options. The FortiOS DHCP server supports up to a maximum of 30 custom options.These optional fields are set in the CLI.

To get to the DHCP server, use the following CLI commands:

```
config system dhcp server
   edit <integer - ID of the specific DHCP server>
```

To configure the options, use the following CLI command:

```
config options
```

Once you are in the options context, create an ID for the table entry, using the following CLI commands:

```
edit <integer>
   set code <integer between 0 - 4294967295 to determine the DHCP option>
   set type [ hex | string | ip ]
   set value <option content for DHCP option types hex and string>
   set ip <option content for DHCP option type ip>
end
```

Excluding addresses in DHCP

If you have a large address range for the DHCP server, you can block a range of addresses that will not be included in the available addresses for the connecting users.

To do this, use the following CLI commands:

```
config system dhcp server
   edit <server_entry_number>
      config exclude-range
         edit <sequence_number>
            set start-ip <address>
            set end-ip <address>
         end
      end
   end
```

Viewing information about DHCP server connections

To view information about DHCP server connections, go to **Monitor > DHCP Monitor**. On this page, you can also add IP addresses to the reserved IP address list.

Breaking an address lease

If you need to end an IP address lease, you can break the lease. This is useful if you have limited addresses and longer lease times, when some leases are no longer necessary, for example, with corporate visitors.

To break a lease, use the following CLI command:

```
execute dhcp lease-clear <ip_address>
```

Interface MTU packet size

You can change the maximum transmission unit (MTU) of the packets that FortiGate transmits to improve network performance. Ideally, the MTU should be the same as the smallest MTU of all the networks between FortiGate and the destination of the packets. If the packets that the FortiGate unit sends are larger than the smallest MTU, they are broken up or fragmented, which slows down transmission. You can easily experiment by lowering the MTU to find an MTU size for optimum network performance.

- 68 to 1500 bytes for static mode
- 576 to 1500 bytes for DHCP mode
- 576 to 1492 bytes for PPPoE mode
- Larger frame sizes (if supported by the FortiGate model), up to 9216 bytes for NP2, NP4, and NP6-accelerated interfaces

This option is available only for physical interfaces. Virtual interfaces associated with a physical interface inherit the physical interface MTU size.

Interfaces on some FortiGate models support frames larger than the traditional 1500 bytes. Jumbo frames are supported on FortiGate models that have either a SOC2 or NP4lite (except for the FortiGate 30D), and on FortiGate 100D series models. For information about your FortiGate model's hardware, see the FortiOS Hardware Acceleration Guide. For other models, contact Fortinet Customer Support for the maximum frame size that is supported.

If you need to send larger frames over a route, all Ethernet devices on that route must support the larger frame size. Otherwise, the larger frames will not be recognized and will be dropped.

If you have standard size and larger size frame traffic on the same interface, routing alone cannot route them to different routes based only on frame size. However, you can use VLANs to make sure the larger frame traffic is routed over network devices that support the larger size. VLANs inherit the MTU size from the parent interface. You must configure the VLAN to include both ends of the route, as well as all switches and routers along the route.

You can configure the MTU packet size. If you select an MTU size larger than your FortiGate model supports, an error message will indicate this. In this situation, try configuring a smaller MTU size until the value is supported.

In transparent mode, if you change the MTU of an interface, you must change the MTU of all interfaces on FortiGate to match the new MTU.

To change the MTU size, use the following CLI commands:

```
config system interface
   edit <interface_name>
      set mtu-override enable
      set mtu <byte_size>
end
```

Interface settings

You configure FortiGate interfaces, both physical and virtual, in **Network > Interfaces** in the FortiGate GUI. There are different options for configuring interfaces when FortiGate is in NAT mode or transparent mode.

On FortiOS Carrier, you can also enable the Gi gatekeeper on each interface for anti-overbilling.

Field	Description
Create New	Select this to add a new interface, zone, or port pair (in transparent mode). Depending on the FortiGate model, you can add a VLAN interface, a loopback interface, an IEEE 802.3ad aggregated interface, or a redundant interface. When VDOMs are enabled, you can also add Inter-VDOM links.

Field	Description
Name	The names of the physical interfaces on FortiGate. This includes any alias names that have been configured. When you combine several interfaces into an aggregate or redundant interface, only the aggregate or redundant interface is listed, and not the component interfaces. If you added VLAN interfaces, they appear in the name list below the physical or aggregated interface to which they have been added. If you added loopback interfaces, they appear in the interface list below the physical interface to which they have been added. If software switch interfaces are configured, you can view them. If your FortiGate model supports AMC modules, the interfaces are named amc-sw1/1, amc-dw1/2, and so on.
Type	The configuration type for the interface.
IP/Netmask	The current IP address and netmask of the interface. In VDOM, when VDOMs are not all in NAT or transparent mode, some values may not be available for display and are displayed as "-".
Access	The administrative access configuration for the interface.
Administrative Status	Indicates if the interface can be accessed for administrative purposes. If the administrative status is a green arrow, an administrator can connect to the interface using the configured access. If the administrative status is a red arrow, the interface is administratively down and cannot be accessed for administrative purposes.
Link Status	The status of the interface physical connection. The link status can be up (green arrow) or down (red arrow). If the link status is up, the interface is connected to the network and accepting traffic. If the link status is down, the interface is either not connected to the network or there is a problem with the connection. You cannot change the link status from the FortiGate GUI, and it typically indicates that an Ethernet cable is plugged into the interface. The link status is only displayed for physical interfaces.
MAC	The MAC address of the interface.
Mode	The addressing mode of the interface. This value can be manual, DHCP, or PPPoE.

Field	Description
Secondary IP	The secondary IP addresses added to the interface.
MTU	The maximum number of bytes per transmission unit for the interface.
Virtual Domain	The virtual domain to which the interface belongs. This column is visible when VDOM configuration is enabled.
VLAN ID	The configured VLAN ID for VLAN subinterfaces.

Interface configuration and settings

To configure an interface, go to **Network > Interfaces**, and select **Create New**.

Name	Enter the name of the interface. Physical interface names cannot be changed.
Alias	Enter an alternate name for a physical interface on the FortiGate unit. This field appears when you edit an existing physical interface. The alias is a maximum of 25 characters. The alias name does not appear in logs.
Link Status	Indicates whether the interface is connected to a network (link status is **Up**) or not (link status is **Down**). This field appears when you edit an existing physical interface.
Type	Select the type of interface you want to add. On some FortiGate models, you can set **Type** to **802.3ad Aggregate** or **Redundant Interface**.
Interface	This is displayed when **Type** is set to **VLAN**. Select the name of the physical interface that you want to add a VLAN interface to. Once created, the VLAN interface is listed below its physical interface in the Interface list. You cannot change the physical interface of a VLAN interface except when you add a new VLAN interface.
VLAN ID	This is displayed when **Type** is set to **VLAN**. Enter the VLAN ID. You cannot change the **VLAN ID** except when you add a new VLAN interface. The VLAN ID must be a number between 1 and 4094. It must match the VLAN ID that the IEEE 802.1Q-compliant router or switch that is connected to the VLAN subinterface adds.

Virtual Domain	Select the virtual domain to add the interface to. Administrator accounts with the super_admin profile can change the **Virtual Domain**.
Physical Interface Members	This section can have two different formats depending on the interface type: • **Software switch interface**: This section is a display-only field that shows the interfaces that belong to the virtual interface of the software switch. • **802.3ad aggregate or Redundant interface**: This section includes the available interface list and the selected interface list. Select interfaces from the **Available Interfaces** list and select the right arrow to add an interface to the **Selected Interface** list.
Addressing mode	Select the addressing mode for the interface: • Select **Manual** and add an **IP/Netmask** for the interface. If IPv6 configuration is enabled, you can add both a IPv4 and an IPv6 IP address. • Select **DHCP** to get the interface IP address and other network settings from a DHCP server. • Select **PPPoE** to get the interface IP address and other network settings from a PPPoE server. • Select **One-Arm Sniffer** to enable the interface as a means to detect possible traffic threats. This option is available on physical ports that are not configured for the primary Internet connection. • Select **Dedicate to FortiAP/FortiSwitch** to have a FortiAP or FortiSwitch device connect exclusively to the interface. This option is available only when you edit a physical interface and it has a static IP address. When you enter the IP address, FortiGate automatically creates a DHCP server using the subnet that you enter. This option is not available on the ADSL interface. The FortiSwitch option is currently available only on the FortiGate 100D.
IP/Netmask	If **Addressing Mode** is set to **Manual**, enter an IPv4 address and subnet mask for the interface. FortiGate interfaces cannot have IP addresses on the same subnet.
IPv6 Address	If **Addressing Mode** is set to **Manual** and IPv6 support is enabled, enter an IPv6 address and subnet mask for the interface. A single interface can have an IPv4 address, IPv6 address, or both.
Administrative Access	Select the types of administrative access that you want to allow for IPv4 connections to this interface.

HTTPS	Allow secure HTTPS connections to the FortiGate GUI through this interface. If configured, this option will enable automatically when you select the **HTTP** option.
PING	The interface responds to pings. Use this setting to verify your installation and for testing.
HTTP	Allow HTTP connections to the FortiGate GUI through this interface. If configured, this option will also enable the **HTTPS** option.
SSH	Allow SSH connections to the CLI through this interface.
SNMP	Allow a remote SNMP manager to request SNMP information by connecting to this interface.
FMG-Access	Allow FortiManager authorization automatically during the communication exchanges between FortiManager and FortiGate devices.
CAPWAP	Allows the FortiGate wireless controller to manage a wireless access point, such as a FortiAP device.
IPv6 Administrative Access	Select the types of administrative access that you want to allow for IPv6 connections to this interface. The types are the same as for Administrative Access.
Security Mode	Select a captive portal for the interface. After you select this, you can define the portal message and the appearance of the GUI that users see when they log into the interface. You can also define one or more user groups that can access the interface.
DHCP Server	Select this to enable a DHCP server for the interface. For more information about configuring a DHCP server on the interface, see DHCP servers and relays.
Device Detection	Select this to allow the interface to be used with BYOD devices, such as iPhones. Define the device definitions by selecting **User & Device > Device Inventory** in the FortiGate GUI.
Enable Explicit Web Proxy	Select this to enable explicit web proxying on this interface. This is available when you enable explicit proxy in the **System Information** Dashboard (**System** > **Dashboard** > **Status**). When you enable this, the interface will be displayed in **System** > **Network** > **Explicit Proxy**, under **Listen on Interfaces**, and web traffic on this interface will be proxied according to the Web Proxy settings. This option is not available for a VLAN interface selection.

Secondary IP Address	Add additional IPv4 addresses to this interface. Select the expand arrow to expand or hide the section.
Comments	Enter a description (up to 63 characters) to describe the interface.
Gi Gatekeeper (FortiOS Carrier only)	For FortiOS Carrier, enable this to enable the Gi firewall as part of the anti-overbilling configuration. You must also configure **Gi Gatekeeper Settings** by selecting **System** > **Admin** > **Settings** in the FortiGate GUI.

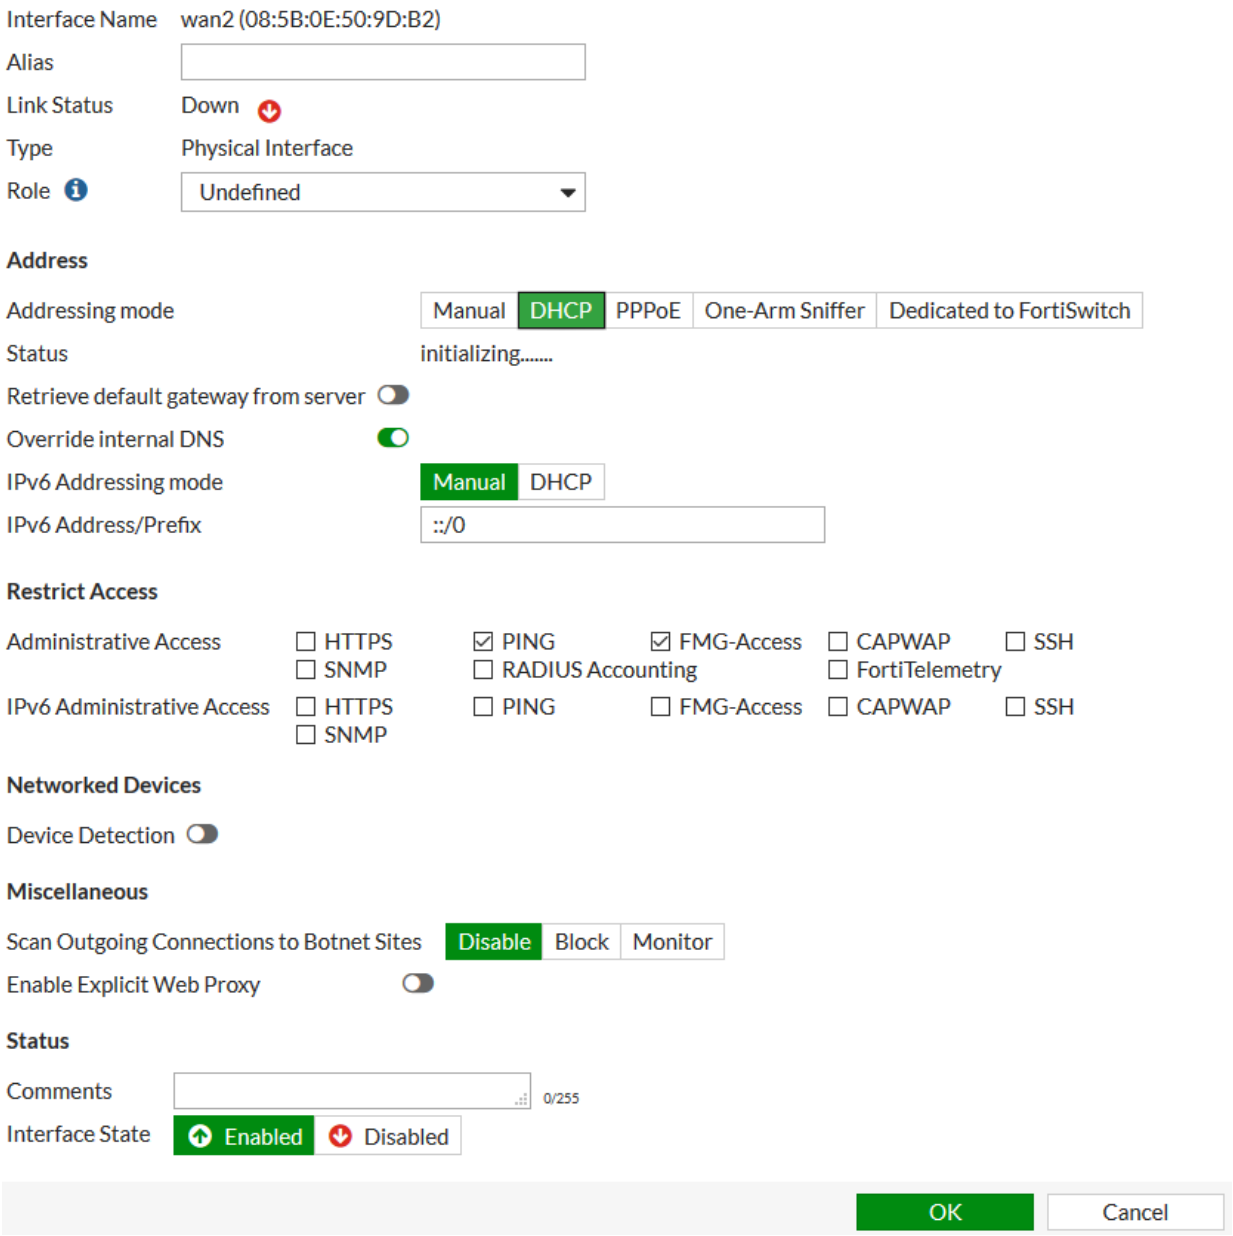

If you assign an interface to be part of a virtual wire pairing, the "role" value is removed from the interface.

Loopback interfaces

A loopback interface is a logical interface that is always up (no physical link dependency) and the attached subnet is always present in the routing table.

The IP address of the FortiGate loopback interface does not depend on one specific external port, and therefore you can access it through several physical or VLAN interfaces. You can configure multiple loopback interfaces in either non-VDOM mode or in each VDOM.

Loopback interfaces still require appropriate firewall policies to allow traffic to and from the interfaces.

A loopback interface can be used with:

- Management access
- BGP (TCP) peering
- PIM RP

Loopback interfaces are a good practice for OSPF. To make troubleshooting OSPF easier, you should set the OSPF router ID the same as the loopback IP address, and remember the management IP addresses (ssh to “router ID”).

You can enable dynamic routing protocols on loopback interfaces.

For black hole static routes, use the black hole route type instead of the loopback interface.

VXLAN loopback binding

A Virtual Extensible LAN (VXLAN) unicast device can bind to a loopback interface as its underlying interface. The IP address of the loopback interface is taken as the source IP address for its outgoing VXLAN packets so the peer knows where to reply. Among the parameters that are passed to the kernel, the ifindex of the loopback interface is not passed down to the kernel, so the kernel can choose the outgoing physical interface. This way, VXLAN traffic can be routed across multiple physical links and it provides resistance to a single point of failure.

You can configure VXLAN loopback binding, using the following CLI commands:

```
config system vxlan
   edit <name>
      set interface <interface>
      set vni <VXLAN network ID>
      set remote-ip <IP address>
   next
end
```

One-armed sniffer

You can use a one-armed sniffer to configure a FortiGate physical interface as a one-arm intrusion detection system (IDS). Traffic sent to the interface is examined for matches to the configured IPS sensor and application control list. Matches are logged and then all received traffic is dropped. Sniffing reports only on attacks. It does not deny or otherwise influence traffic.

You can use the one-arm sniffer to configure FortiGate to operate as an IDS appliance by sniffing network traffic for attacks without actually processing the packets. To configure one-arm IDS, you enable sniffer mode on a

FortiGate interface and connect the interface to a hub, or to the SPAN port of a switch that is processing network traffic.

To assign an interface as a sniffer interface, select **Network > Interfaces**, edit the interface and select **One-Arm Sniffer**.

If the check box is not available, it means the interface is in use. Ensure that the interface is not selected in any firewall policies, routes, virtual IPs, or other features in which a physical interface is specified.

Field	Description
Enable Filters	Select this to include filters that define a more granular sniff of network traffic. Select specific hosts, ports, VLANs, and protocols. In all cases, enter a number or number range for the filtering type. For protocol values, the standard protocols are: • UDP - 17 • TCP - 6 • ICMP - 1
Include IPv6 Packets	If your network is running both IPv4 and IPv6 addressing, select this to sniff both addressing types. Otherwise, FortiGate will sniff only IPv4 traffic.
Include Non-IP Packets	Select this for a more intense scan of content in the traffic.
Security Profiles	IPS sensors and application control lists allow you to select specific sensors and applications that you want to identify within the traffic.

Interface Name wan2 (08:5B:0E:50:9D:B2)
Alias
Link Status Down
Type Physical Interface
Role Undefined

Address

Addressing mode Manual | DHCP | PPPoE | One-Arm Sniffer | Dedicated to FortiSwitch
☑ Enable Filters
Host(s)
Port(s)
VLAN(s)
Protocol
☐ Include IPv6 Packets
☐ Include Non-IP Packets

Security Profiles

Enable AntiVirus Edit Sniffer Profile
Enable Web Filter Edit Sniffer Profile
Enable Application Control Edit Sniffer Profile
Enable CASI Profile Edit Sniffer Profile
Enable IPS Edit Sniffer Profile

Logging Options

Log Allowed Traffic Security Events | All Sessions
Scan Outgoing Connections to Botnet Sites Disable | Block | Monitor

Ports preassigned as sniffer ports

Some FortiGate models have ports preconfigured as sniffer ports, by default. The models and ports preconfigered in sniffer mode are as follows:

- FortiGate 300D
  - Port4
  - Port8
- FortiGate 500D
  - Port5
  - Port6
  - Port13
  - Port14

Physical ports

FortiGate has several physical ports that you can connect Ethernet or optical cables to. Depending on the FortiGate model, it can have between 4 and 40 physical ports. Some units have a grouping of ports labeled as lan, that provide built-in switch functionality.

The port names, as labeled on FortiGate, appear in the FortiGate GUI in the **Unit Operation** widget on the dashboard. They also appear when you configure the interfaces, in **Network > Interfaces**.

You can hover over the ports to see information about each port, such as the name of the port and the IP address.

For example, the following diagram shows the 22 interfaces of the FortiGate 100 D (Generation 2) as they appear in the dashboard in the FortiGate GUI.

Two of the physical ports on the FortiGate 100D (Generation 2) are SFP ports. These ports share the numbers 15 and 16 with RJ-45 ports. Because of this, when SFP port 15 is used, RJ-45 port 15 cannot be used, and vice versa. These ports also share the same MAC address.

Configuring the FortiGate 100D ports

Normally, you can configure the internal interface as a single interface that is shared by all physical interface connections (a switch). The switch mode feature has two states: switch mode and interface mode. Switch mode is the default mode, with only one interface and one address for the entire internal switch. Interface mode allows you to configure each of the physical interface connections of the internal switch separately. This allows you to assign different subnets and netmasks to each of the internal physical interface connections.

The larger FortiGate models may also include Advanced Mezzanine Cards (AMC), which can provide additional interfaces (Ethernet or optical, with throughput enhancements for more efficient handling of specialized traffic. These interfaces appear in FortiOS as port amc/sw1, amc/sw2, and so on.

Displaying information about the status of transceivers

You can display information about the status of transceivers installed in FortiGate SFP/SFP+ interfaces, in the FortiGate CLI.

The `get system interface transceiver` command lists all of the SFP/SFP+ interfaces on FortiGate. If the interfaces include transceivers, the command output displays information about them, such as the vendor name, part number, and serial number. It also includes details about transceiver operation, such as temperature, voltage, and optical transmission power, which you can use to diagnose transmission problems.

The following example shows an output from using this command:

```
get system interface transceiver
...
Interface port14 - Transceiver is not detected.
Interface port15 - SFP/SFP+
  Vendor Name  :            FIBERXON INC.
  Part No.     :            FTM-8012C-SLG
  Serial No.   :            101680071708917
Interface port16 - SFP/SFP+
  Vendor Name  :            FINISAR CORP.
  Part No.     :            FCLF-8521-3
  Serial No.   :            PS62ENQ

                                      Optical      Optical      Optical
SFP/SFP+     Temperature  Voltage     Tx Bias      Tx Power     Rx Power
Interface    (Celsius)    (Volts)     (mA)         (dBm)        (dBm)
------------ ------------ ------------ ------------ ------------ ------------
port15        N/A          N/A          N/A          N/A          N/A
port16        N/A          N/A          N/A          N/A          N/A
  ++ : high alarm, + : high warning, - : low warning, -- : low alarm, ? : suspect.
```

You can use this command on most FortiGate models that have SFP/SFP+ interfaces.

Split port support

The 5001D 40 GB can be split into 4 10 GB ports.You can do this through a combination of hardware and software configuration. You use a specific 40 GB connector to connect to the 40 GB port and typically, the other end of the fibre optic cable connects to another 40 GB port. However, you can use a special cable that is a single 40 GB connector at one end and 4 10 GB connections at the other end. To use this setup, you also have to configure the port to be a split port.

You can configure this, using the following CLI commands:

```
config system global
   set port-split port1 port2
   end
```

The ports will be checked to make sure that they are not in use or referenced by other policy configurations. If they are in use, the command is aborted. Changing the port to be a split port requires a system reboot.

PPPoE addressing mode on an interface

If you configure the interface to use PPPoE, the FortiGate unit automatically broadcasts a PPPoE request from the interface.

The FortiGate units support many PPPoE RFC features (RFC 2516) including unnumbered IPs, initial discovery timeout and PPPoE Active Discovery Terminate (PADT).

PPPoE is only configurable in the web-based manager on desktop FortiGate units. 1U FortiGates and up must be configured in the CLI using the commands:

```
config system interface
   edit <port_name>
      set mode pppoe
      set username <ISP_username>
```

```
        set password <ISP_password>
        set idle-timeout <seconds>
        set distance <integer>
        set ipunnumbered <unumbered-IP>
        set disc-retry-timeout <seconds>
        set padt-retry-timeout <seconds>
        set lcp-echo-interval <seconds>
        set dns-server-override {enable | disable}
end
```

Configure PPPoE on an interface in **Network > Interfaces**. The following table describes the PPPoE status information when PPPoE is configured for an interface.

Field	Description
Status	Displays PPPoE status messages as the FortiGate unit connects to the PPPoE server and gets addressing information. Select Status to refresh the addressing mode status message. The status is only displayed if you selected **Edit**. Status can be any one of the following 4 messages.
Initializing	No activity.
Connecting	The interface is attempting to connect to the PPPoE server.
Connected	The interface retrieves an IP address, netmask, and other settings from the PPPoE server. When the status is connected, PPPoE connection information is displayed.
Failed	The interface was unable to retrieve an IP address and other information from the PPPoE server.
Reconnect	Select to reconnect to the PPPoE server. Only displayed if Status is connected.
User Name	The user name for the PPPoE account.
Password	The password for the PPPoE account.
Unnumbered IP	Specify the IP address for the interface. If your ISP has assigned you a block of IP addresses, use one of them. Otherwise, this IP address can be the same as the IP address of another interface or can be any IP address.
Initial Disc Timeout	Enter Initial discovery timeout. Enter the time to wait before starting to retry a PPPoE discovery.

Field	Description
Initial PADT timeout	Enter Initial PPPoE Active Discovery Terminate (PADT) timeout, in seconds. Use this timeout to shut down the PPPoE session if it is idle for the specified number of seconds. PADT must be supported by your ISP. Set the Initial PADT timeout to 0 to disable.
Distance	Enter the administrative distance for the default gateway retrieved from the PPPoE server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route. The default distance for the default gateway is 1.
Retrieve default gateway from server	Enable to retrieve a default gateway IP address from a PPPoE server. The default gateway is added to the static routing table.
Override internal DNS	Enable to replace the DNS server IP addresses on the System DNS page with the DNS addresses retrieved from the PPPoE server. When VDOMs are enabled, you can override the internal DNS only on the management VDOM.

Probing interfaces

Server probes can be used on interfaces. In order for this to occur, the probe response mode must first be configured, then the probe response must be allowed administrative access on the interface. The probe response mode can one of the following:

Mode	Description
`none`	Disable probe
`http-probe`	HTTP probe
`twamp`	Two-Way Active Measurement Protocol

Both steps must be done through the CLI.

Configuring the probe

```
config system probe-response
   set mode http-probe
end
```

Allowing the probe response to have administrative access to the interface

```
config system interface
   edit <port>
      set allowaccess probe-response
   end
```

Enhanced TWAMP Light functionality with server/controller functionality

TWAMP(Two-Way Active Measurement Protocol) Light is a simplified architecture within the TWAMP standard. Its purpose is to measure the round trip IP performance between any two devices within a network that supports the protocol. FortiOS operates in more than just the role of responder/reflector.The server/controller functionality has been added.

Redundant interfaces

On some models, you can combine two or more physical interfaces to provide link redundancy. This feature allows you to connect to two or more switches to ensure connectivity if one physical interface, or the equipment on that interface, fails.

In a redundant interface, traffic travels only over one interface at a time. This differs from an aggregated interface where traffic travels over all interfaces for distribution of increased bandwidth. This difference means that redundant interfaces can have more robust configurations with fewer possible points of failure. This is important in a fully-meshed HA configuration.

An interface can be in a redundant interface if:

- It is a physical interface, not a VLAN interface
- It is not already part of an aggregated or redundant interface
- It is in the same VDOM as the redundant interface
- It has no defined IP address
- It is not configured for DHCP or PPPoE
- It has no DHCP server or relay configured on it
- It does not have any VLAN subinterfaces
- It is not referenced in any security policy, VIP, or multicast policy
- It is not monitored by HA
- It is not one of the FortiGate-5000 series backplane interfaces

When an interface is included in a redundant interface, it is not listed on the **Network > Interfaces** page. You cannot configure the interface individually and it is not available for inclusion in security policies, VIPs, or routing.

Dual Internet connections

Dual internet connections, also referred to as dual WAN or redundant Internet connections, refers to using two FortiGate interfaces to connect to the Internet. Dual Internet connections can be used in three ways:

- Redundant interfaces: if one interface goes down, the second interface automatically becomes the main Internet connection
- Load sharing: to ensure better throughput
- Combination of redundancy and load sharing

Redundant interfaces

Redundant interfaces ensure that if your Internet access is no longer available through a certain port, the FortiGate unit will use an alternate port to connect to the Internet.

Configuring redundant interfaces

In this scenario, two interfaces, WAN1 and WAN2, are connected to the Internet using two different ISPs. WAN1 is the primary connection. In the event of a failure of WAN1, WAN2 automatically becomes the connection to the Internet. For this configuration to function correctly, you need to configure three settings:

- Configure a link health monitor to determine when the primary interface (WAN1) is down and when the connection returns
- Configure a default route for each interface
- Configure security policies to allow traffic through each interface to the internal network

Link health monitor

Adding a link health monitor is required for routing failover traffic. A link health monitor will confirm the connectivity of the device's interface.

To add a link health monitor

```
config system link-monitor
   edit "Example1"
      set srcint <Interface_sending_probe>
      set server <ISP_IP_address>
      set protocol <Ping or http>
      set gateway-ip <the_gateway_IP_to_reach_the_server_if_required>
      set failtime <failure_count>
      set interval <seconds>
      set update-cascade-interface enable
      set update-static-route enable
      set status enable
end
```

Routing

You need to configure a default route for each interface and indicate which route is preferred by specifying the distance. The lower distance is declared active and placed higher in the routing table.

When you have dual WAN interfaces that are configured to provide failover, you might not be able to connect to the backup WAN interface because the FortiGate unit may not route traffic (even responses) out of the backup interface. The FortiGate unit performs a reverse path lookup to prevent spoofed traffic. If an entry cannot be found in the routing table that sends the return traffic out the same interface, the incoming traffic is dropped.

To configure the routing of the two interfaces - web-based manager

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	For an IPv4 route, enter a subnet of `0.0.0.0/0.0.0.0` For an IPv6 route, enter a subnet of `::/0`
Gateway	Enter the gateway address
Interface	Select the primary connection. For example, **WAN1**.
Administrative Distance	Leave as the default of `10`.

3. Repeat these steps to set **Interface** to **WAN2** and **Administrative Distance** to `20`.

To configure the IPv4 routing of the two interfaces - CLI

```
config router static
   edit 0
      set dst 0.0.0.0 0.0.0.0
      set device WAN1
      set gateway <gateway_address>
      set distance 10
   next
   edit 0
      set dst 0.0.0.0 0.0.0.0
      set device WAN2
      set gateway <gateway_address>
      set distance 20
   next
end
```

To configure the IPv6 routing of the two interfaces - CLI

```
config router static6
   edit 0
      set dst ::/0
      set device WAN1
      set gateway <gateway_address>
      set distance 10
   next
   edit 0
      set dst ::/0
      set device WAN2
      set gateway <gateway_address>
      set distance 20
   next
end
```

Security policies

When creating security policies, you need to configure duplicate policies to ensure that after traffic fails over WAN1, regular traffic will be allowed to pass through WAN2 as it did with WAN1. This ensures that failover will occur with minimal effect to users. For more information about creating security policies, see the Firewall Guide.

Load sharing

Load sharing allows you to use both connections to the Internet at the same time, but does not provide failover support. When configuring load sharing, you need to make sure that routing is configured for both external ports (for example, WAN1 and WAN2) have static routes with the same distance and priority.

For more information about load sharing, see the Advanced Routing Guide.

Link redundancy and load sharing

In this scenario, both links are available to distribute Internet traffic over both links. Should one of the interfaces fail, the FortiGate unit will continue to send traffic over the other active interface. Configuration is similar to the Redundant interfaces configuration, with the main difference being that the configured routes should have equal distance settings.

This means both routes will remain active in the routing table. To make one interface the preferred interface, use a default policy route to indicate the interface that is preferred for accessing the Internet. If traffic matches the security policy, the policy overrides all entries in the routing table, including connected routes. You may need to add specific policy routes that override these default policy routes.

To redirect traffic over the secondary interface, create policy routes to direct some traffic onto it rather than the primary interface. When adding the policy route, only define the outgoing interface and leave the gateway blank. This ensures that the policy route will not be active when the link is down.

SSL VPN and WAN link load balancing

You can set virtual WAN link interfaces as destination interfaces in firewall policies for WAN link load balancing, when SSL VPN is the source interface. For example, you can log in to FortiGate using an SSL VPN for traffic inspection and then have outbound traffic load balanced by WAN link load balancing.

You can set a virtual WAN link interface as a destination interface in a firewall policy where SSL VPN is the source interface, using either the FortiGate GUI (FortiOS 5.6.1 and later) or CLI.

In the CLI, use the following CLI commands:

```
config firewall policy
   edit <policy_id>
      set dstintf virtual-wan-link
end
```

Secondary IP addresses to an interface

If an interface is configured with a manual or static IP address, you can also add secondary static IP addresses to the interface. Adding secondary IP addresses effectively adds multiple IP addresses to the interface. Secondary IP addresses cannot be assigned using DCHP or PPPoE.

All of the IP addresses added to an interface are associated with the single MAC address of the physical interface , and all secondary IP addresses are in the same VDOM as the interface that they are added to. You configure interface status detection for gateway load balancing separately for each secondary IP addresses. As with all other interface IP addresses, secondary IP addresses cannot be on the same subnet as any other primary or secondary IP address assigned to a FortiGate interface unless they are in separate VDOMs.

To configure a secondary IP address, go to **Network > Interfaces**, select **Edit** or **Create New** and select the **Secondary IP Address** check box.

Software switch

A software switch, or soft switch, is a virtual switch that is implemented at the software, or firmware level, rather than the hardware level. A software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch, you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiGate unit.

It can also be useful if you require more hardware ports for the switch on a FortiGate unit. For example, if your FortiGate unit has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces such as those with FortiWiFi and FortiAP unit.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to each interface are not regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

There are a few things to consider when setting up a software switch:

- Ensure you create a backup of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiGate unit. If you accidentally combine too many ports, you will need a way to undo any errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiGate unit. For example, DHCP servers, security policies, and so on.
- For increased security, you can create a captive portal for the switch, allowing only specific user groups access to the resources connected to the switch.
- To add an interface to a software switch, the interface cannot be referenced by the existing configuration. It must also have its IP address set to 0.0.0.0/0.0.0.0.

To create a software switch - CLI

```
config system switch-interface
   edit <switch-name>
      set type switch
      set member <interface_list>
end
config system interface
   edit <switch_name>
      set ip <ip_address>
      set allowaccess https ssh ping
end
```

Soft switch example

For this example, the wireless interface (Wi-Fi) needs to be on the same subnet as the DMZ1 interface to facilitate wireless syncing from an iPhone and a local computer. The synching between two subnets is problematic. By putting both interfaces on the same subnet, the synching will work. The software switch will accomplish this.

In this example, the soft switch includes a wireless interface. Remember to configure any wireless security before proceeding. If you leave this interface open without any password or other security, it leaves open access to not only the wireless interface but to any other interfaces and devices connected within the software switch.

Clear the interfaces and back up the configuration

First, ensure that the interfaces are not being used with any other security policy or other use on the FortiGate unit. Check the Wi-Fi and DMZ1 ports to ensure that DHCP is not enabled on the interface and there are no other dependencies with these interfaces.

Next, save the current configuration, in the event something does not work, recovery can be quick.

Merge the interfaces

The plan is to merge the Wi-Fi port and DMZ1 port. This will create a software switch with a name of “synchro” with an IP address of 10.10.21.12. The following steps will create the switch, add the IP address and set administrative access for HTTPS, SSH, and Ping.

To merge the interfaces - CLI

```
config system switch-interface
   edit synchro
      set type switch
      set member dmz1 wifi
end
config system interface
   edit synchro
      set ip 10.10.21.12
      set allowaccess https ssh ping
end
```

Final steps

With the switch set up, you can add security policies, DHCP servers, and any other configuration that you would normally do to configure interfaces on the FortiGate unit.

Virtual switch

Virtual switch feature enables you create virtual switches on top of the physical switch(es) with designated interfaces/ports so that a virtual switch can build up its forwarding table through learning and forward traffic accordingly. When traffic is forwarded among interfaces belonging to the same virtual switch, the traffic does not need to go up to the software stack, but is forwarded directly by the switch. When traffic has to be relayed to

interfaces not on the virtual switch, the traffic will go through the normal data path and be offloaded to NP4, when possible.

This feature is only available on mid- to high-end FortiGate units, including the 100D, 600C, 1000C, and 1240B.

To enable and configure the virtual switch, enter the following CLI commands:

```
config system virtual-switch
   edit vs1
      set physical-switch sw0
         config port
            edit 1
               set port port1
               set speed xx
               set duplex xx
               set status [up|down]
            edit 2
               set port port2
               set ...
            end
      end
end
```

Support for 802.1x fallback and 802.1x dynamic VLANs

There are four modes when enabling 802.1x on a virtual switch interface:

Mode	Description
Default	In this mode, it works as it did previously.
Fallback	In fallback mode, the virtual switch will be treated as a master. Only one slave can refer to a fallback master. Those ports in the master virtual switch are always authorized. After passing 802.1x authentication, the ports will be stay authorized and moved to its slave virtual switch.
Dynamic-vlan	In dynamic-vlan mode, the virtual switch will also be treated as a master. However, many slaves can refer to a dynamic-vlan master. Those ports in the master virtual switch are always un-authorized. After passing 802.1x/MAB authentication, the ports will be set to authorized and moved to one of its slave virtual switches.
Slave	In slave mode, a master must be set through security-8021x-master attribute. A slave virtual switch will use its master virtual switch's security-groups settings for authentication.

CLI example for fallback mode:

```
config system virtual-switch
   edit "fallsw"
      set physical-switch "sw0"
      config port
   end
   edit "trust"
      set physical-switch "sw0"
   end
```

```
config system interface
   edit "fallsw"
      set vdom "root"
      set ip 192.168.20.1 255.255.255.0
      set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
          proberesponse capwap
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode fallback(fallback mode master switch)
      set security-groups "rds-grp"(the usergroup for 802.1x)
      set snmp-index 10
      next
   edit "trust"
      set vdom "root"
      set ip 192.168.22.1 255.255.255.0
      set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
          proberesponse
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode slave(slave mode switch)
      set security-8021x-master "fallsw" (its master switch)
      set snmp-index 6
      next
   end
```

CLI example for dynamic-vlan mode:

```
config system virtual-switch
   edit "internal"
      set physical-switch "sw0"
   edit "lan-trust"
      set physical-switch "sw0"
      next
   edit "lan-vlan1000"
      set physical-switch "sw0"
      next
   edit "lan-vlan2000"
      set physical-switch "sw0"
      config port
      edit "internal1" (normally we should not add port in slave switch. This is used if
      user wants to manually add one port in slave)
      end
   end
config system interface
   edit "internal"
      set vdom "root"
      set ip 192.168.11.99 255.255.255.0
      set allowaccess ping https ssh http fgfm capwap
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode dynamic-vlan<------dynamic-vlan mode master switch
      set security-groups "rds-grp"<------the usergroup for 802.1x
      set snmp-index 15
      next
   edit "lan-trust"
      set vdom "root"
      set ip 192.168.111.99 255.255.255.0
```

```
      set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
          proberesponse capwap
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode slave<-----slave mode switch
      set security-8021x-master "internal"<-----its master switch
      set snmp-index 7
      next
   edit "lan-vlan1000"
      set vdom "root"
      set ip 192.168.110.1 255.255.255.0
      set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
          proberesponse capwap
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode slave<-----slave mode switch
      set security-8021x-master "internal"<-----its master switch
      set security-8021x-dynamic-vlan-id 1000 <-----the matching vlan id for this virtual
      switch
      set snmp-index 16
      next
   edit "lan-vlan2000"
      set vdom "root"
      set ip 192.168.220.1 255.255.255.0
      set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
          proberesponse
      capwap
      set type hard-switch
      set security-mode 802.1X
      set security-8021x-mode slave
      set security-8021x-master "internal"
      set security-8021x-dynamic-vlan-id 2000
      set snmp-index 17
      end
   config user group
      edit "rds-grp"
         set dynamic-vlan-id 4000(default vlan id if there is no vlan attribute return
             from server)
         set member "190"
      end
```

Zones

Zones are a group of one or more FortiGate interfaces, both physical and virtual, that you can apply security policies to control inbound and outbound traffic. Grouping interfaces and VLAN subinterfaces into zones simplifies the creation of security policies where a number of network segments can use the same policy settings and protection profiles. When you add a zone, you select the names of the interfaces and VLAN subinterfaces to add to the zone. Each interface still has its own address and routing is still done between interfaces, that is, routing is not affected by zones. Security policies can also be created to control the flow of intra-zone traffic.

For example, the network includes three separate groups of users representing different entities on the company network. While each group has its own set of port and VLANs, in each area, they can all use the same security policy and protection profiles to access the Internet. Rather than the administrator making nine separate security

policies, the administrator can add the required interfaces to a zone and create three policies, making administration simpler.

You can configure policies for connections to and from a zone, but not between interfaces in a zone.

The following example shows how to set up a zone to include the internal interface and a VLAN.

To create a zone - web-based manager

1. Go to **Network > Interfaces**.
2. Select the arrow on the **Create New** button and select **Zone**.
3. Enter a zone name of `Zone_1`.
4. Select the required **Interface Members**.
5. Select **OK**.

To create a zone - CLI

```
config system zone
   edit Zone_1
      set interface internal VLAN_1
end
```

Virtual domains

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. A single FortiGate unit is then flexible enough to serve multiple departments of an organization, separate organizations, or to act as the basis for a service provider's managed security service.

VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations. By default, each FortiGate unit has a VDOM named root. This VDOM includes all of the FortiGate physical interfaces, modem, virtual LAN (VLAN) subinterfaces, zones, security policies, routing settings, and VPN settings.

When a packet enters a VDOM, it is confined to that VDOM. In a VDOM, you can create security policies for connections between VLAN subinterfaces or zones in the VDOM. Packets do not cross the virtual domain border internally. To travel between VDOMs, a packet must pass through a firewall on a physical interface. The packet then arrives at another VDOM on a different interface, but it must pass through another firewall before entering the VDOM. Both VDOMs are on the same FortiGate unit. Inter-VDOMs change this behavior because they are internal interfaces; however, their packets go through all the same security measures as on physical interfaces.

The following example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First, enable VDOMs on the FortiGate unit. When you enable VDOMs, the FortiGate unit will log you out.

For desktop and low-end FortiGate units, you use the CLI to enable VDOMs. Once you enable VDOMs, all further configuration can be done using the web-based manager or the CLI. On larger FortiGate units, you can use the web-based manager or the CLI to enable VDOMs.

To enable VDOMs - web-based manager

1. Go to **Dashboard**.
2. In the **System Information** widget, select **Enable** for **Virtual Domain**.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all VDOMs:

To enable VDOMs - CLI

```
config system global
   set vdom-admin enable
end
```

Next, add the VDOM called accounting.

To add a VDOM - web-based manager

1. Go to **System > VDOM**, and select **Create New**.
2. Enter the VDOM name `accounting`.
3. Select **OK**.

To add a VDOM - CLI

```
config vdom
   edit <new_vdom_name>
end
```

With the VDOM created, you can assign a physical interface to it, and assign it an IP address.

To assign physical interface to the accounting Virtual Domain - web-based manager

1. Go to **Network > Interfaces**.
2. Select the DMZ2 port row and select **Edit**.
3. For the **Virtual Domain** drop-down list, select **accounting**.
4. Select the **Addressing Mode** of **Manual**.
5. Enter the IP address for the port of 10.13.101.100/24.
6. Set the **Administrative Access** to **HTTPS** and **SSH**.
7. Select **OK**.

To assign physical interface to the accounting Virtual Domain - CLI

```
config global
   config system interface
      edit dmz2
         set vdom accounting
         set ip 10.13.101.100/24
         set allowaccess https ssh
   next
end
```

Wireless

A wireless interface is similar to a physical interface except it does not include a physical connection. The FortiWiFi units allow you to add multiple wireless interfaces that can be available at the same time. On FortiWiFi

units, you can configure the device to be either an access point or a wireless client. As an access point, the FortiWiFi unit can have separate SSIDs, each on their own subnet for wireless access. In client mode, the FortiWiFi has only one SSID, and is used as a receiver to allow remote users to connect to the existing network using wireless protocols.

Wireless interfaces also require additional security measures to ensure the session does not get hijacked and data tampered or stolen.

For more information about configuring wireless interfaces see the *Deploying Wireless Networks Guide*.

VLANs

Virtual Local Area Networks (VLANs) multiply the capabilities of your FortiGate unit, and can also provide added network security. VLANs use ID tags to logically separate devices on a network into smaller broadcast domains. These smaller domains forward packets only to devices that are part of that VLAN domain. This reduces traffic and increases network security. The IEEE 802.1Q standard defines VLANs. All layer-2 and layer-3 devices along a route must be 802.1Q-compliant to support VLANs along that route.

A Local Area Network (LAN) is a group of connected computers and devices that are arranged into network broadcast domains. A LAN broadcast domain includes all the computers that receive a packet broadcast from any computer in that broadcast domain. A switch will automatically forward the packets to all of its ports. In contrast, routers do not automatically forward network broadcast packets. This means routers separate broadcast domains. If a network has only switches and no routers, that network is considered one broadcast domain, no matter how large or small it is. Smaller broadcast domains are more efficient because fewer devices receive unnecessary packets. They are more secure as well because a hacker reading traffic on the network will have access to only a small portion of the network instead of the entire network's traffic.

VLANs reduce the size of the broadcast domains by only forwarding packets to interfaces that are part of that VLAN or part of a VLAN trunk link. Trunk links form switch-to-switch or switch-to-router connections, and forward traffic for all VLANs. This enables a VLAN to include devices that are part of the same broadcast domain, but physically distant from each other.

VLAN ID tags consist of a 4-byte frame extension that switches and routers apply to every packet sent and received in the VLAN. Workstations and desktop computers, which are commonly originators or destinations of network traffic, are not an active part of the VLAN process. All the VLAN tagging and tag removal is done after the packet has left the computer.

Any FortiGate unit without VDOMs enabled can have a maximum of 255 interfaces in transparent operating mode. The same is true for any single VDOM. In NAT mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in transparent operating mode, you need to configure multiple VDOMs that enable you to divide the total number of interfaces over all the VDOMs.

One example of an application of VLANs is a company's accounting department. Accounting computers may be located at both main and branch offices. However, accounting computers need to communicate with each other frequently and require increased security. VLANs allow the accounting network traffic to be sent only to accounting computers and to connect accounting computers in different locations as if they were on the same physical subnet.

This guide uses the term "packet" to refer to both layer-2 frames and layer-3 packets.

On a layer-2 switch, you can have only one VLAN subinterface per physical interface, unless that interface is configured as a trunk link. Trunk links can transport traffic for multiple VLANs to other parts of the network.

On a FortiGate unit, you can add multiple VLANs to the same physical interface. However, VLAN subinterfaces added to the same physical interface cannot have the same VLAN ID or have IP addresses on the same subnet. You can add VLAN subinterfaces with the same VLAN ID to different physical interfaces.

Creating VLAN subinterfaces with the same VLAN ID does not create an internal connection between them. For example a VLAN ID of 300 on port1 and VLAN ID of 300 on port2 are allowed, but they are not connected. Their relationship is the same as between any two FortiGate network interfaces.

FortiGate unit interfaces cannot have overlapping IP addresses, the IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask. This rule helps prevent a broadcast storm or other similar network problems.

The following example shows how to add a VLAN, called vlan_accounting, on the FortiGate unit internal interface with an IP address of 10.13.101.101.

To add a VLAN - web-based manager

1. Go to **Network > Interfaces.**
2. Select **Create New** and click on **Interface**.
   The **Type** is set to VLAN, by default.
3. Enter a name for the **VLAN** to `vlan_accounting`.
4. Select the **Internal** interface.
5. Enter the **VLAN ID**.
   The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together.
6. Select the **Addressing Mode** of **Manual**.
7. Enter the IP address for the port of 10.13.101.101/24.
8. Set the **Administrative Access** to **HTTPS** and **SSH**.
9. Select **OK**.

To add a VLAN - CLI

```
config system interface
   edit VLAN_1
      set interface internal
      set type vlan
      set vlanid 100
      set ip 10.13.101.101/24
      set allowaccess https ssh
   next
end
```

VLANs in NAT mode

In NAT mode the FortiGate unit functions as a layer-3 device. In this mode, the FortiGate unit controls the flow of packets between VLANs, but can also remove VLAN tags from incoming VLAN packets. The FortiGate unit can also forward untagged packets to other networks, such as the Internet.

In NAT mode, the FortiGate unit supports VLAN trunk links with IEEE 802.1Q-compliant switches, or routers. The trunk link transports VLAN-tagged packets between physical subnets or networks. When you add VLAN sub-interfaces to the FortiGate unit physical interfaces, the VLANs have IDs that match the VLAN IDs of packets on the trunk link. The FortiGate unit directs packets with VLAN IDs to sub-interfaces with matching IDs.

You can define VLAN sub-interfaces on all FortiGate physical interfaces. However, if multiple virtual domains are configured on the FortiGate unit, you will have access to only the physical interfaces on your virtual domain. The FortiGate unit can tag packets leaving on a VLAN subinterface. It can also remove VLAN tags from incoming packets and add a different VLAN tag to outgoing packets.

Normally in VLAN configurations, the FortiGate unit's internal interface is connected to a VLAN trunk, and the external interface connects to an Internet router that is not configured for VLANs. In this configuration the FortiGate unit can apply different policies for traffic on each VLAN interface connected to the internal interface, which results in less network traffic and better security.

Adding VLAN subinterfaces

A VLAN subinterface, also called a VLAN, is a virtual interface on a physical interface. The subinterface allows routing of VLAN tagged packets using that physical interface, but it is separate from any other traffic on the physical interface.

Adding a VLAN subinterface includes configuring:

- Physical interface
- IP address and netmask
- VLAN ID
- VDOM

Physical interface

The term VLAN subinterface correctly implies the VLAN interface is not a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network router that is configured for this VLAN. Without that router, the VLAN will not be connected to the network, and VLAN traffic will not be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

When you are working with interfaces on your FortiGate unit, use the **Column Settings** on the Interface display to make sure the information you need is displayed. When working with VLANs, it is useful to position the **VLAN ID** column close to the IP address. If you are working with VDOMs, including the **Virtual Domain** column as well will help you troubleshoot problems more quickly.

To view the Interface display, go to **Network > Interfaces**.

IP address and netmask

FortiGate unit interfaces cannot have overlapping IP addresses. The IP addresses of all interfaces must be on different subnets. This rule applies to both physical and virtual interfaces, such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask pair. This rule helps prevent a broadcast storm or other similar network problems.

If you are unable to change your existing configurations to prevent IP overlap, enter the CLI command `config system settings` and `set allow-subnet-overlap enable` to allow IP address overlap. If you enter this command, multiple VLAN interfaces can have an IP address that is part of a subnet used by another interface. This command is recommended for advanced users only.

VLAN ID

The VLAN ID is part of the VLAN tag added to the packets by VLAN switches and routers. The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together. VLAN ID 0 is used only for high priority frames, and 4095 is reserved.

All devices along a route must support the VLAN ID of the traffic along that route. Otherwise, the traffic will be discarded before reaching its destination. For example, if your computer is part of VLAN_100 and a co-worker on a different floor of your building is also on the same VLAN_100, you can communicate with each other over VLAN_100, only if all the switches and routers support VLANs and are configured to pass along VLAN_100 traffic properly. Otherwise, any traffic you send to your co-worker will be blocked or will not be delivered.

VDOM

If VDOMs are enabled, each VLAN subinterface must belong to a VDOM. This rule also applies for physical interfaces.

Interface-related CLI commands require a VDOM to be specified, regardless of whether the FortiGate unit has VDOMs enabled.

VLAN subinterfaces on separate VDOMs cannot communicate directly with each other. In this situation, the VLAN traffic must exit the FortiGate unit and re-enter the unit, passing through firewalls in both directions. This situation is the same for physical interfaces.

A VLAN subinterface can belong to a different VDOM than the physical interface it is part of. This is because the traffic on the VLAN is handled separately from the other traffic on that interface. This is one of the main strengths of VLANs.

The following procedure will add a VLAN subinterface called `VLAN_100` to the FortiGate internal interface with a VLAN ID of 100. It will have an IP address and netmask of `172.100.1.1/255.255.255.0`, and allow HTTPS and PING administrative access. Note that in the CLI, you must enter "`set type vlan`" before setting the vlanid, and that the allowaccess protocols are lower case.

To add a VLAN subinterface in NAT mode - web-based manager

1. If **Current VDOM** appears at the bottom left of the screen, select **Global** from the list of VDOMs.
2. Go to **Network > Interfaces**.
3. Select **Create New** to add a VLAN subinterface.
4. Enter the following:

VLAN Name	VLAN_100
Type	VLAN
Interface	internal
VLAN ID	100
Addressing Mod	Manual
IP/Netmask	172.100.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET

5. Select **OK**.

To view the new VLAN subinterface, select the expand arrow next to the parent physical interface (the internal interface). This will expand the display to show all VLAN subinterfaces on this physical interface. If there is no expand arrow displayed, there are no subinterfaces configured on that physical interface.

For each VLAN, the list displays the name of the VLAN, and, depending on column settings, its IP address, the Administrative access you selected for it, the VLAN ID number, and which VDOM it belongs to if VDOMs are enabled.

To add a VLAN subinterface in NAT mode - CLI

```
config system interface
   edit VLAN_100
      set interface internal
      set type vlan
      set vlanid 100
      set ip 172.100.1.1 255.255.255.0
      set allowaccess https ping
end
```

Configuring security policies and routing

Once you have created a VLAN subinterface on the FortiGate unit, you need to configure security policies and routing for that VLAN. Without these, the FortiGate unit will not pass VLAN traffic to its intended destination. Security policies direct traffic through the FortiGate unit between interfaces. Routing directs traffic across the network.

Configuring security policies

Security policies permit communication between the FortiGate unit's network interfaces based on source and destination IP addresses. Interfaces that communicate with the VLAN interface need security policies to permit traffic to pass between them and the VLAN interface.

Each VLAN needs a security policy for each of the following connections the VLAN will be using:

- From this VLAN to an external network
- From an external network to this VLAN
- From this VLAN to another VLAN in the same virtual domain on the FortiGate unit
- From another VLAN to this VLAN in the same virtual domain on the FortiGate unit

The packets on each VLAN are subject to antivirus scans and other security profiles measures as they pass through the FortiGate unit.

Configuring routing

As a minimum, you need to configure a default static route to a gateway with access to an external network for outbound packets. In more complex cases, you must configure different static or dynamic routes based on packet source and destination addresses.

As with firewalls, you must configure routes for VLAN traffic. VLANs need routing and a gateway configured to send and receive packets outside their local subnet just as physical interfaces do. The type of routing you configure, static or dynamic, will depend on the routing used by the subnet and interfaces you are connecting to. Dynamic routing can be routing information protocol (RIP), border gateway protocol (BGP), open shortest path first (OSPF), or multicast.

If you enable SSH, PING, HTTPS and HTTP on the VLAN, you can use those protocols to troubleshoot your routing and test that it is properly configured. Enabling logging on the interfaces and using CLI diagnose commands, such as `diagnose sniff packet <interface_name>`, can also help locate any possible configuration or hardware issues.

VLANs in transparent mode

In transparent mode, the FortiGate unit behaves like a layer-2 bridge but can still provide services such as antivirus scanning, web filtering, spam filtering and intrusion protection to traffic. There are some limitations in transparent mode because you cannot use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in transparent mode apply to IEEE 802.1Q VLAN trunks passing through the unit.

VLANs and transparent mode

You can insert the FortiGate unit operating in transparent mode into the VLAN trunk without making changes to your network. In a typical configuration, the FortiGate unit internal interface accepts VLAN packets on a VLAN trunk from a VLAN switch or router connected to internal network VLANs. The FortiGate external interface forwards VLAN-tagged packets through another VLAN trunk to an external VLAN switch or router and on to external networks such as the Internet. You can configure the unit to apply different policies for traffic on each VLAN in the trunk.

To pass VLAN traffic through the FortiGate unit, add two VLAN subinterfaces with the same VLAN ID, one to the internal interface and the other to the external interface. You then create a security policy to permit packets to flow from the internal VLAN interface to the external VLAN interface. If required, you create another security policy to permit packets to flow from the external VLAN interface to the internal VLAN interface. Typically in transparent mode, you do not permit packets to move between different VLANs. Network protection features, such as spam filtering, web filtering, and anti-virus scanning, are applied through the Security Profiles specified in each security policy, enabling very detailed control over traffic.

When the FortiGate unit receives a VLAN-tagged packet at a physical interface, it directs the packet to the VLAN subinterface with the matching VLAN ID. The VLAN tag is removed from the packet, and the FortiGate unit then applies security policies using the same method it uses for non-VLAN packets. If the packet exits the FortiGate unit through a VLAN subinterface, the VLAN ID for that subinterface is added to the packet and the packet is sent to the corresponding physical interface.

General configuration steps

There are two essential steps to configure your FortiGate unit to work with VLANs in transparent mode:

- Add VLAN subinterfaces
- Create security policies

You can also configure the Security Profiles that manage antivirus scanning, web filtering and spam filtering. For more information about Security Profiles, see the Security Profiles Guide.

Add VLAN subinterfaces

The VLAN ID of each VLAN subinterface must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch. The VLAN ID can be any number between 1 and 4094, with 0 being used only for high priority frames and 4095 being reserved. You add VLAN subinterfaces to the physical interface that receives VLAN-tagged packets.

For this example, we are creating a VLAN called internal_v225 on the internal interface, with a VLAN ID of 225. Administrative access is enabled for HTTPS and SSH. VDOMs are not enabled.

To add VLAN subinterfaces in transparent mode - web-based manager

1. Go to **Network > Interfaces**.
2. Select **Create New** and click on **Interfaces**.
3. Enter the following information and select **OK**.

Name	internal_v225
Type	VLAN
Interface	internal
VLAN ID	225
Administrative Access	Enable HTTPS, and SSH. These are very secure access methods.

The FortiGate unit adds the new subinterface to the interface that you selected.

Repeat steps 2 and 3 to add additional VLANs. You will need to change the **VLAN ID**, **Name**, and possibly **Interface** when adding additional VLANs.

To add VLAN subinterfaces in transparent mode - CLI

```
config system interface
   edit internal_v225
      set interface internal
      set vlanid 225
      set allowaccess HTTPS SSH
      set description "VLAN 225 on internal interface"
      set vdom root
end
```

Create security policies

In transparent mode, the FortiGate unit performs antivirus and antispam scanning on each VLAN's packets as they pass through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another.

To add security policies for VLAN subinterfaces - web based manager

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** to add firewall addresses that match the source and destination IP addresses of VLAN packets.
3. Go to **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6Policy** and select **Create New**.
4. From the **Incoming Interface/Zone** list, select the VLAN interface where packets enter the unit.
5. From the **Outgoing Interface/Zone** list, select the VLAN interface where packets exit the unit.
6. Select the **Source** and **Destination** Address names that you added in step 2.
7. Select **OK**.

To add security policies for VLAN subinterfaces - CLI

```
config firewall address
   edit incoming_VLAN_address
      set associated-interface <incoming_VLAN_interface>
      set type ipmask
      set subnet <IPv4_address_mask)
   next
   edit outgoing_VLAN_address
      set associated-interface <outgoing_VLAN_interface>
      set type ipmask
      set subnet <IPv4_address_mask>
   next
end
config firewall policy or config firewall policy6
   edit <unused_policy_number>
      set srcintf <incoming_VLAN_interface>
      set srcaddr incoming_VLAN_address
      set destintf <outgoing_VLAN_interface>
      set destaddr outgoing_VLAN_address
      set schedule always
      set service <protocol_to_allow_on VLAN>
      set action ACCEPT
   next
end
```

VLAN switching and routing

VLAN switching takes place on the Open Systems Interconnect (OSI) model layer-2, just like other network switching. VLAN routing takes place on the OSI model layer-3. The difference between them is that during VLAN switching, VLAN packets are simply forwarded to their destination. This is different from VLAN routing where devices can open the VLAN packets and change their VLAN ID tags to route the packets to a new destination.

VLAN layer-2 switching

Ethernet switches are layer-2 devices, and generally are 802.1Q compliant. Layer 2 refers to the second layer of the seven layer OSI basic networking model, called the Data Link layer. FortiGate units act as layer-2 switches or bridges when they are in transparent mode. The units simply tag and forward the VLAN traffic or receive and remove the tags from the packets. A layer-2 device does not inspect incoming packets or change their contents; it only adds or removes tags and routes the packet.

A VLAN can have any number of physical interfaces assigned to it. Multiple VLANs can be assigned to the same physical interface. Typically two or more physical interfaces are assigned to a VLAN, one for incoming and one for outgoing traffic. Multiple VLANs can be configured on one FortiGate unit, including trunk links.

VLAN layer-3 routing

Routers are layer-3 devices. Layer 3 refers to the third layer of the OSI networking model, the Network layer. FortiGate units in NAT mode act as layer-3 devices. As with layer 2, FortiGate units acting as layer-3 devices are 802.1Q-compliant.

The main difference between layer-2 and layer-3 devices is how they process VLAN tags. Layer-2 switches just add, read and remove the tags. They do not alter the tags or do any other high-level actions. Layer-3 routers not only add, read and remove tags but also analyze the data frame and its contents. This analysis allows layer-3 routers to change the VLAN tag if it is appropriate and send the data frame out on a different VLAN.

In a layer-3 environment, the 802.1Q-compliant router receives the data frame and assigns a VLAN ID. The router then forwards the data frame to other members of the same VLAN broadcast domain. The broadcast domain can include local ports, layer-2 devices and layer-3 devices such as routers and firewalls. When a layer-3 device receives the data frame, the device removes the VLAN tag and examines its contents to decide what to do with the data frame. The layer-3 device considers:

- Source and destination addresses
- Protocol
- Port number

The data frame may be forwarded to another VLAN, sent to a regular non-VLAN-tagged network or just forwarded to the same VLAN as a layer-2 switch would do. Or, the data frame may be discarded if the proper security policy has been configured to do so.

Layer-2 and ARP traffic

By default, FortiGate units do not pass layer-2 traffic. If there are layer-2 protocols such as IPX, PPTP or L2TP in use on your network, you need to configure your FortiGate unit interfaces to pass these protocols without blocking. Another type of layer-2 traffic is Address Resolution Protocol (ARP) traffic.

You can allow these layer-2 protocols using the CLI command:

```
config system interface
   edit <name_str>
      set l2forward enable
   end
```

where `<name_str>` is the name of an interface.

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem as follows:

```
config vdom
   edit <vdom_name>
      config system interface
      edit <name_str>
      set l2forward enable
   end
end
```

If you enable layer-2 traffic, you may experience a problem if packets are allowed to repeatedly loop through the network. This repeated looping, very similar to a broadcast storm, occurs when you have more than one layer-2 path to a destination. Traffic may overflow and bring your network to a halt. You can break the loop by enabling Spanning Tree Protocol (STP) on your network's switches and routers.

STP forwarding

The FortiGate unit does not participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, reducing available bandwidth to nothing.

If you use your FortiGate unit in a network topology that relies on STP for network loop protection, you need to make changes to your FortiGate configuration. Otherwise, STP recognizes your FortiGate unit as a blocked link and forwards the data to another path. By default, your FortiGate unit blocks STP as well as other non-IP protocol traffic.

Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the external interface:

```
config system interface
   edit external
      set l2forward enable
      set stpforward enable
end
```

By substituting different commands for `stpforward enable`, you can also allow layer-2 protocols such as IPX, PPTP or L2TP to be used on the network.

STP support for FortiGate models with hardware switches

STP (Spanning Tree Protocol) used to be available only on the old style switch mode for the internal ports. You can now activate STP on the hardware switches found in the newer FortiGate models. These models use a virtual switch to simulate the old switch Mode for the internal ports.

You can enable STP, using the following CLI commands:

```
config system interface
   edit lan
   set stp [enable | disable]
   end
```

ARP traffic

Address Resolution Protocol (ARP) packets are vital to communication on a network and ARP support is enabled on FortiGate unit interfaces, by default. Normally you want ARP packets to pass through the FortiGate unit, especially if it is sitting between a client and a server or between a client and a router.

ARP traffic can cause problems, especially in transparent mode where ARP packets arriving on one interface are sent to all other interfaces including VLAN subinterfaces. Some layer-2 switches become unstable when they detect the same MAC address originating on more than one switch interface or from more than one VLAN. This instability can occur if the layer-2 switch does not maintain separate MAC address tables for each VLAN. Unstable switches may reset and cause network traffic to slow down considerably.

The default ARP timeout value is 5 minutes (300 seconds). So usually ARP entries are removed after 5 minutes. However, some conditions can cause ARP entries to remain on the list for a longer time. This is not a configurable value. To view the ARP list, enter the `get system arp` CLI command.

Proxy ARP extensions

You can extend the proxy ARP configuration to an IP address range instead of a single IP address. When you configure `proxy-arp`, in addition to setting the IP address, you can also set the `end-ip` address. If you do not set this, the proxy ARP will be a single address, as before. The following is an example CLI configuration, using the new setting:

```
config system proxy-arp
   edit 1
      set interface "internal"
      set ip 192.168.1.100
      set end-ip 192.168.1.102
      end
```

Multiple VDOMs solution

By default, physical interfaces are in the root domain. If you do not configure any of your VLANs in the root VDOM, it will not matter how many interfaces are in the root VDOM.

The multiple VDOMs solution is to configure multiple VDOMs on the FortiGate unit, one for each VLAN. In this solution, you configure one inbound and one outbound VLAN interface in each VDOM. ARP packets are not forwarded between VDOMs. This configuration limits the VLANs in a VDOM and correspondingly reduces the administration needed per VDOM.

As a result of this configuration, the switches do not receive multiple ARP packets with duplicate MACs. Instead, the switches receive ARP packets with different VLAN IDs and different MACs. Your switches are stable.

However, you should **not** use the multiple VDOMs solution under any of the following conditions:

- You have more VLANs than licensed VDOMs
- You do not have enough physical interfaces

Instead, use one of two possible solutions, depending on which operation mode you are using:

- In NAT mode, you can use the `vlanforward` CLI command.
- In transparent mode, you can use the `forward-domain` CLI command. But you still need to be careful in some rare configurations.

Vlanforward solution

If you are using NAT mode, the solution is to use the `vlanforward` CLI command for the interface in question. By default, this command is enabled and will forward VLAN traffic to all VLANs on this interface. When disabled, each VLAN on this physical interface can send traffic only to the same VLAN. There is no cross-talk between VLANs, and ARP packets are forced to take one path along the network which prevents the multiple paths problem.

In the following example, `vlanforward` is disabled on port1. All VLANs configured on port1 will be separate and will not forward any traffic to each other.

```
config system interface
```

```
    edit port1
       set vlanforward disable
end
```

Forward-domain solution

If you are using transparent mode, the solution is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic. It is like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0. The many benefits of this solution include reduced administration, the need for fewer physical interfaces, and the availability of more flexible network solutions.

In the following example, forward-domain collision group 340 includes VLAN 340 traffic on port1 and untagged traffic on port 2. Forward-domain collision group 341 includes VLAN 341 traffic on port 1 and untagged traffic on port 3. All other interfaces are part of forward-domain collision group 0 by default. This configuration separates VLANs 340 and 341 from each other on port 1.

Use these CLI commands:

```
config system interface
   edit port2
      set forward_domain 340
   next
   edit port3
      set forward_domain 341
   next
   edit port1-340
      set forward_domain 340
      set interface port1
      set vlanid 340
   next
   edit port1-341
      set forward_domain 341
      set interface port1
      set vlanid 341
end
```

You may experience connection issues with layer-2 traffic, such as ping, if your network configuration has:

- Packets going through the FortiGate unit in transparent mode more than once
- More than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled

Now IPS and AV is applied the first time packets go through the FortiGate unit, but not on subsequent passes. Applying IPS and AV only to this first pass fixes the network layer-2 related connection issues.

Asymmetric routing

You might discover unexpectedly that hosts on some networks are unable to reach certain other networks. This occurs when request and response packets follow different paths. If the FortiGate unit recognizes the response packets, but not the requests, it blocks the packets as invalid. Also, if the FortiGate unit recognizes the same packets repeated on multiple interfaces, it blocks the session as a potential attack.

This is asymmetric routing. By default, the FortiGate unit blocks packets or drops the session when this happens. You can configure the FortiGate unit to permit asymmetric routing by using the following CLI commands:

```
config system settings
   set asymroute enable
end
```

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem as follows:

```
config vdom
   edit <vdom_name>
      config system settings
      set asymroute enable
   end
end
```

If this solves your blocked traffic issue, you know that asymmetric routing is the cause. But allowing asymmetric routing is not the best solution, because it reduces the security of your network.

For a long-term solution, it is better to change your routing configuration or change how your FortiGate unit connects to your network.

If you enable asymmetric routing, antivirus and intrusion prevention systems will not be effective. Your FortiGate unit will be unaware of connections and treat each packet individually. It will become a stateless firewall.

Configuring IPv4 and IPv6 ICMP traffic inspection

In order for the inspection of asymmetric ICMP traffic not to affect TCP and UDP traffic, you can enable or disable ICMP traffic inspection for traffic being routed asymmetrically for both IPv4 and IPv6.

To configure ICMP traffic inspeciton, use the following CLI commands:

- IPv4:

```
config system settings
   set asymroute-icmp
   end
```

- IPv6:

```
config system settings
   set asymroute6-icmp
   end
```

NetBIOS

Computers running Microsoft Windows operating systems that are connected through a network rely on a WINS server to resolve host names to IP addresses. The hosts communicate with the WINS server by using the NetBIOS protocol.

To support this type of network, you need to enable the forwarding of NetBIOS requests to a WINS server. The following example will forward NetBIOS requests on the internal interface for the WINS server located at an IP address of 192.168.111.222.

```
config system interface
   edit internal
      set netbios_forward enable
      set wins-ip 192.168.111.222
end
```

These commands apply only in NAT mode. If VDOMs are enabled, these commands are per VDOM. You must set them for each VDOM that has the problem.

Too many VLAN interfaces

Any virtual domain can have a maximum of 255 interfaces in transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. NAT mode supports from 255 to 8192 depending on the FortiGate model. This total number of interfaces includes VLANs, other virtual interfaces, and physical interfaces.

Your FortiGate unit may allow you to configure more interfaces than this. However, if you configure more than 255 interfaces, your system will become unstable and, over time, will not work properly. As all interfaces are used, they will overflow the routing table that stores the interface information, and connections will fail. When you try to add more interfaces, an error message will state that the maximum limit has already been reached.

If you see this error message, chances are you already have too many VLANs on your system and your routing has become unstable. To verify, delete a VLAN and try to add it back. If you have too many, you will not be able to add it back on to the system. In this case, you will need to remove enough interfaces (including VLANs) so that the total number of interfaces drops to 255 or less. After doing this, you should also reboot your FortiGate unit to clean up its memory and buffers, or you will continue to experience unstable behavior.

To configure more than 255 interfaces on your FortiGate unit in transparent mode, you have to configure multiple VDOMs, each with many VLANs. However, if you want to create more than the default 10 VDOMs (or a maximum of 2550 interfaces), you must buy a license for additional VDOMs and your FortiGate must be able to be licensed for more than 10 VDOMs.

With these extra licenses, you can configure up to 500 VDOMs, with each VDOM containing up to 255 VLANs in transparent mode. This is a theoretical maximum of over 127 500 interfaces. However, system resources will quickly get used up before reaching that theoretical maximum. To achieve the maximum number of VDOMs, you need to have top-end hardware with the most resources possible.

In NAT mode, if you have a top-end model, the maximum interfaces per VDOM can be as high as 8192, which is enough for all the VLANs in your configuration.

Your FortiGate unit has limited resources, such as CPU load and memory, that are divided between all configured VDOMs. When running 250 or more VDOMs, you may need to monitor the system resources to ensure there is enough to support the configured traffic processing.

Troubleshooting VLAN issues

Several problems can occur with your VLANs. Since VLANs are interfaces with IP addresses, they behave as interfaces and can have similar problems that you can diagnose with tools, such as ping, traceroute, packet sniffing, and diag debug.

Botnet and command-and-control protection

You can configure botnet and command-and-control traffic protection, in the FortiGate GUI or CLI.

In the GUI, you can use select the **Scan Outgoing Connections to Botnet Sites** option on the **Interfaces** page. The options are **Disable**, **Block**, and **Monitor**.

In the CLI, you can configure the botnet scan on the interface, using the following commands:

```
config system interface
   edit <interface>
      set scan-botnet-connections [disable | block | monitor]
      end
```

You can also enable the scanning of botnet and command-and-control traffic in the following policies:

- Firewall policies:

```
config firewall policy
   edit <policyid>
      set scan-botnet-connections [disable | block | monitor]
      end
```

- Firewall explicit proxy policies:

```
config firewall explicit-proxy-policy
   edit <policyid>
      set scan-botnet-connections [disable | block | monitor]
      end
```

- Firewall interface policy:

```
config firewall interface-policy
   edit <policyid>
      set scan-botnet-connections [disable | block | monitor]
      end
```

- Firewall sniffer:

```
config firewall sniffer
   edit <policyid>
      set scan-botnet-connections [disable | block | monitor]
      end
```

DNS

A Domain Name System (DNS) server is a public service that converts symbolic node names to IP addresses. A DNS server implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with their computer IP addresses. This allows you to use readable locations, such as fortinet.com, when you browse the Internet. FortiOS supports DNS configuration for both IPv4 and IPv6 addressing.

FortiGate includes default DNS server addresses. However, you should change these addresses to ones that your Internet Service Provider (ISP) provides. The defaults are DNS proxies and are not as reliable as those from your ISP.

Within FortiOS, there are two DNS configuration options. Each option provides a specific service and both options can work together to provide a complete DNS solution.

DNS settings

You configure basic DNS queries on interfaces that connect to the Internet. When a user requests a website, FortiGate looks to the configured DNS servers to provide the IP address of the website in order to know which server to contact to complete the transaction.

You configure DNS server addresses by selecting **Network > DNS**, and then specifying the DNS server addresses. These addresses are typically supplied by your ISP. If you have local Microsoft domains on the network, you can enter a domain name in the **Local Domain Name** field.

In a situation where all three fields are configured, FortiGate first looks to the local domain. If no match is found, FortiGate sends a request to the external DNS servers.

If virtual domains (VDOM) are enabled, you create a DNS database in each VDOM. All of the interfaces in a VDOM share the DNS database in that VDOM.

Additional DNS CLI configuration

Additional DNS configuration options are available in the CLI, using the `config system dns` command. Within this command, you can also set the following commands:

Command	Description
`dns-cache-limit`	Set how many DNS entries are stored in the cache. Entries that remain in the cache provide a quicker response to requests than going out to the Internet to get the same information.
`dns-cache-ttl`	Set how long entries remain in the cache, in seconds. Possible values are 60 to 86400 (default is 24 hours).

Command	Description
`cache-notfound-responses`	When you enable this, any DNS requests that are returned with NOT FOUND can be stored in the cache.
`source-ip`	Define a dedicated IP address for communications with the DNS server.

DDNS

If your ISP changes your external IP address on a regular basis, and you have a static domain name, you can configure the external interface to use a dynamic DNS (DDNS) service. This ensures that external users and customers can always connect to your company firewall. If you have a FortiGuard subscription, you can use FortiGuard as the DDNS server.

You can configure FortiGuard as the DDNS server, in the FortiGate GUI or CLI.

To configure FortiGuard as the DDNS server in the FortiGate GUI, select **Network > DNS** and enable **FortiGuard DDNS**. Then select the interface with the dynamic connection, which DDNS server you have an account with, your domain name, and account information. If your DDNS server is not on the list, there is a generic option where you can provide your DDNS server information.

To configure FortiGuard as the DDNS server in the FortiGate CLI, use the following CLI commands:

```
config system fortiguard
   set ddns-server-ip
   set ddns-server-port
end
```

If you do not have a FortiGuard subscription or want to use a different DDNS server, you can configure DDNS in the CLI. You can configure a DDNS for each interface. Only the first configured port appears in the FortiGate GUI. Additional commands vary depending on the DDNS server you select. Use the following CLI commands:

```
config system ddns
   edit <DDNS_ID>
      set monitor-interface <external_interface>
      set ddns-server <ddns_server_selection>
end
```

Configuring FortiGate to refresh DDNS IP addresses

You can configure FortiGate to refresh DDNS IP addresses. FortiGate periodically checks the DDNS server that is configured. Use the following CLI commands:

```
config system ddns
   edit <1>
      set ddns-server FortiGuardDDNS
      set use-public-ip enable
      set update-interval seconds
   end
```

The possible values for update-interval are 60 to 2592000 seconds, and the default is 300 seconds.

TLS support for DDNS updates

When cleartext is disabled, FortiGate uses the SSL connection to send and receive Dynamic DNS services (DDNS) updates.

To disable cleartext, use the following CLI commands:

```
config system ddns
   set clear-text disable
   end
```

The ssl-certificate name can also be set in the same location using the command:

```
set ssl-certificate <cert_name>
```

DDNS update override for DHCP

DHCP server has an override command option that allows DHCP server communications to go through DDNS to perform updates for the DHCP client. This enforces a DDS update of the AA field every time, even if the DHCP client does not request it. This allows the support of the allow/ignore/deny client-updates options.

You can enable DDNS update override, using the following CLI commands:

```
config system dhcp server
   edit <0>
      set ddns-update_override enable
      end
```

FortiDDNS registration to a public IP address

Fortinet's Dynamic DNS services (FortiDDNS) can be registered to a public IP address even if the FortiGate model does not have any physical interfaces on the Internet. This applies to when FortiGate is behind other networking devices that are employing NAT. You can configure this in the GUI and the CLI.

DNS servers

You can also create local DNS servers for your network. Depending on your requirements, you can manually maintain your entries (master DNS server) or use it as a jumping point, where the server refers to an outside source (slave DNS server). A local master DNS server works similarly to the DNS server addresses configured in **Network > DNS**, but you must manually add all entries. This allows you to add a local DNS server to include specific URL and IP address combinations.

The DNS server options are not visible in the FortiGate GUI, by default. To enable the server, select **System > Feature Visibility**, select **DNS Database**, and select **Apply**.

While a master DNS server is an easy method to include regularly used addresses to save on going to an outside DNS server, it is not recommended to make it the authoritative DNS server. IP addresses may change and maintaining any type of list can become labor-intensive.

It is best to use a FortiGate master DNS server for local services. For example, a company has a web server in their DMZ that internal users (employees) and external users (customers or remote employees) access. When internal users access a website, a request for the website is sent out to the DNS server on the Internet, which

then returns an IP address or virtual IP address. After the company configures an internal DNS server, the same website request is resolved internally to the internal web server IP address. This minimizes inbound and outbound traffic, and access time.

As a slave DNS server, FortiGate refers to an external or alternate source as a way to obtain the URL and IP address combination. This is useful if there is a master DNS server for a large company, where a list is maintained. Satellite offices can then connect to the master DNS server to obtain the correct addressing.

The DNS server entries do not allow CNAME entries, as per RFC 1912, section 2.4.

Configure a master DNS server - web-based manager

1. Select **Network > DNS Servers**, and select **Create New** for **DNS Database**.
2. Select the **Type** of **Master**.
3. Select the **View** as **Shadow**.
4. The view is the accessibility of the DNS server. Selecting **Public**, external users can access, or use, the DNS server. Selecting **Shadow**, only internal users can use it.
5. Enter the DNS **Zone**, for example, `WebServer`.
6. Enter the domain name for the zone, for example `example.com`.
7. Enter the hostname of the DNS server, for example, `Corporate`.
8. Enter the contact address for the administrator, for example, `admin@example.com`.
9. Set **Authoritative** to **Disable**.
10. Select **OK**.
11. Enter the DNS entries for the server by selecting **Create New**.
12. Select the **Type**, for example, **Address (A)**.
13. Enter the **Hostname**, for example `web.example.com`.
14. Enter the remaining information, which varies depending on the **Type** selected.
15. Select **OK**.

Configure a master DNS server - CLI

```
config system dns-database
   edit WebServer
      set domain example.com
      set type master
      set view shadow
      set ttl 86400
      set primary-name corporate
      set contact admin@exmple.com
      set authoritative disable
         config dns-entry
            edit 1
               set hostname web.example.com
               set type A
               set ip 192.168.21.12
               set status enable
         end
      end
```

Configuring a recursive DNS

You can set an option to ensure this type of DNS server is not the authoritative server. When configured, the FortiGate unit will check its internal DNS server (master or slave). If the request cannot be fulfilled, it will look to the external DNS servers. This is known as a split DNS configuration.

You can also have FortiGate look to an internal server if the master or slave does not fulfill the request, using the following CLI commands:

```
config system dns-database
   edit example.com
      ...
      set view shadow
end
```

For this behavior to work completely, you must set the DNS query for the external interface to be recursive.

Configure a recursive DNS - web-based manager

1. Go to **Network > DNS Servers**, and select **Create New** for **DNS Service on Interface**.
2. Select the **Interface**.
3. Select the **Mode** to **Recursive**.
4. Select **OK**.

Configure a recursive DNS - CLI

```
config system dns-server
   edit wan1
      set mode recursive
end
```

Configuring IPv6 Router Advertisement options for DNS configuration

FortiGate supports the following RFC 6106 IPv6 Router Advertisement options:

- Obtaining DNS search list options from upstream DHCPv6 servers
- Sending the DNS search list through Router Advertisement
- Sending the DNS search list through the FortiGate DHCP server
- Sending DNS search list option to downstream clients with Router Advertisements that use a static prefix (FortiOS version 5.6.1 and later)
- Sending recursive DNS server option to downstream clients with Router Advertisements that use a static prefix (FortiOS version 5.6.1 and later)

Obtain the DNS search list options from upstream DHCPv6 servers - CLI

```
config system interface
   edit wan1
      config ipv6
         set dhcp6-prefix-delegation enable
      next
   next
end
```

Send DNS search lists through Router Advertisement - CLI

```
config system interface
   edit port 1
      config IPv6
         set ip6-address 2001:10::/64
         set ip6-mode static
         set ip6-send-adv enable
         config ip6-delegated-prefix-list
         edit 1
            set upstream-interface WAN
            set subnet 0:0:0:11::/64
            set autonomous-flag enable
            set onlink-flag enable
         next
      next
   end
end
```

Send the DNS search lists through the FortiGate DHCP server - CLI

You can use the `dns-search-list delegated` command to send DNS search list option to downstream clients with Router Advertisements that use a static prefix, using the following CLI commands:

```
config system dhcp6 server
   edit 1
      set interface port2
      set upstream-interface WAN
      set ip-mode delegated
      set dns-service delegated
      set dns-search-list delegated
      set subnet 0:0:0:12::/64
   next
end
```

Send DNS search list option to downstream clients with Router Advertisements that use a static prefix - CLI

In FortiOS 5.6.1 and later, you can use the `set dnssl <DNS search list option>` command to send DNS search list option to downstream clients with Router Advertisements that use a static prefix, using the following CLI commands:

```
config system interface
   edit port1
      config ipv6
         config ip6-prefix-list
            edit <2001:db8::/64>
                  set autonomous-flag enable
                  set onlink-flag enable
                  set rdnss 2001:1470:8000::66 2001:1470:8000::72
                  set dnssl <DNS search list option>
               end
```

Send recursive DNS server option to downstream clients with Router Advertisements that use a static prefix - CLI

In FortiOS 5.6.1 and later, you can use the `set rdnss <recursive DNS search option>` command to send Recursive DNS server option to downstream clients with Router Advertisements that use a static prefix, using the following CLI commands:

```
config system interface
   edit port1
      config ipv6
         config ip6-prefix-list
            edit <2001:db8::/64>
                  set autonomous-flag enable
                  set onlink-flag enable
                  set rdnss 2001:1470:8000::66 2001:1470:8000::72
                  set dnssl <DNS search list option>
               end
```

Viewing the Internet Service Database

The Internet Service Database in the FortiGate GUI contains detailed information about services that are available on the Internet, such as DNS servers that Adobe, Google, Fortinet, and Apple provide. For each service, the database shows the IP addresses of the servers that host the service, and the port and protocol number that each IP address uses.To view the Internet Service Database, select **Policy & Objects > Internet Service Database** in the FortiGate GUI.

Advanced static routing

Advanced static routing includes features and concepts that are used in more complex networks.

Routing concepts

Many routing concepts apply to static routing. However without first understanding these basic concepts, it is difficult to understand the more complex dynamic routing.

Routing in VDOMs

Routing on FortiGate units is configured per-VDOM. This means if VDOMs are enabled, you must enter a VDOM to do any routing configuration. This allows each VDOM to operate independently, with its own default routes and routing configuration.

In this guide, the procedures assume your FortiGate unit has VDOMs disabled. This is stated in the assumptions for the examples. If you have VDOMs enabled, you will need to perform the following steps in addition to the procedure's steps.

To route in VDOMs - web-based manager

Select the VDOM that you want to view or configure at the bottom of the main menu.

To route in VDOMs - CLI

Before following any CLI routing procedures with VDOMs enabled, enter the following commands. For this example, it is assumed you will be working in the root VDOM. Change root to the name of your selected VDOM as needed.

```
config vdom
   edit root
```

Following these commands, you can enter any routing CLI commands, as normal.

Default route

The default route is used if either there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

All routers, including FortiGate units, are shipped with default routes in place. This allows customers to set up and become operational more quickly. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

Adding a static route

1. To add or edit a static route, go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	Enter the destination IP address and netmask. A value of `0.0.0.0/0.0.0.0` is universal.
Gateway	Enter the gateway IP address.
Interface	Select the name of the interface that the static route will connect through.
Administrative Distance	Enter the distance value, which will affect which routes are selected first by different protocols for route management or load balancing. The default is 10.
Advanced Options	Optionally, expand **Advanced Options** and enter a **Priority**, which will artificially weight the route during route selection. The higher the priority number, the less likely the route is to be selected over other routes. The default is 0.

Enabling or disabling individual static routes

You can enable or disable individual static routes.

To configure IPv4 static routes, use the following CLI commands:

```
config route static
   edit 0
      set status [enable|disable]
   end
```

To configure IPv6 static routes, use the following CLI commands:

```
config route static6
   edit 0
      set status [enable|disable]
      end
```

Configuring FQDNs as a destination address in static routes

You can configure FQDN firewall addresses as destination addresses in a static route, using either the GUI or the CLI.

In the GUI, to add an FQDN firewall address to a static route in the firewall address configuration, enable the **Static Route Configuration** option. Then, when configuring the static route, set **Destination** to **Named Address**.

In the CLI, use the following CLI commands:

First, configure the firewall FQDN address:

```
config firewall address
   edit 'Fortinet-Documentation-Website'
      set type fqdn
```

```
        set fqdn docs.fortinet.com
        set allow-routing enable
    end
```

Next, add the FQDN address to a static route.

```
config router static
    edit 0
        set dstaddr Fortinet-Documentation-Website
        ...
    end
```

Routing table

When two computers are directly connected, there is no need for routing because each computer knows exactly where to find the other computer. They communicate directly.

Networking computers allows many computers to communicate with each other. This requires each computer to have an IP address to identify its location to the other computers. This is much like a mailing address, where you will not receive your postal mail at home if you do not have an address for people to send mail to. The routing table on a computer is much like an address book used to mail letters to people, where the routing table maintains a list of how to reach computers. Routing tables may also include information about the quality of service (QoS) of the route, and the interface associated with the route if the device has multiple interfaces.

Looking at routing as delivering letters is more simple than reality. In reality, routers lose power or have bad cabling, network equipment is moved without warning, and other such events happen that prevent static routes from reaching their destinations. When any changes, such as these, happen along a static route, traffic can no longer reach the destination and the route goes down. Dynamic routing can address these changes to ensure that traffic still reaches its destination. The process of realizing there is a problem, backtracking, and finding a route that is operational, is called convergence. If there is fast convergence in a network, users will not even know that re-routing is taking place.

The routing table for any device on the network has a limited size. For this reason, routes that are not used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes, which are the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information but it also looks up the source information to ensure that it exists. If there is no source to be found, that packet is dropped because the router assumes it is an error or an attack on the network.

The routing table is used to store routes that are learned. The routing table for any device on the network has a limited size. For this reason, routes that are not used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes, which are the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Viewing the routing table

You can view the routing table in the FortiGate GUI. By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of "any/all" packets.

To display the routes in the routing table, go to **Monitor > Routing Monitor**. Select **Static & Dynamic** to view the routes.

You can also monitor policy routes. Select **Policy** to list the active policy routes on the FortiGate and see information about them. The active policy routes include policy routes that you create, SD-WAN rules, and Internet service static routes. It also supports downstream devices in the Security Fabric.

The following figure show an example of the static and dynamic routes in the Routing Monitor list.

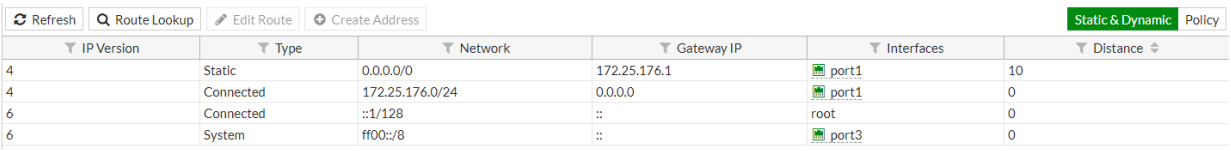

IP Version	Type	Network	Gateway IP	Interfaces	Distance
4	Static	0.0.0.0/0	172.25.176.1	port1	10
4	Connected	172.25.176.0/24	0.0.0.0	port1	0
6	Connected	::1/128	::	root	0
6	System	ff00::/8	::	port3	0

The following figure show an example of the policy routes in the Routing Monitor list.

IP Version	From	Source	To	Destination	Gateway IP	Protocol	Action
4	mgmt1	10.10.10.0/255.255.255.0	Any	0.0.0.0/0.0.0.0	172.20.121.2	Any	Route

Field	Description
IP Version	Shows whether the route is IPv4 or IPv6. IPv6 routes are displayed only if IPv6 is enabled in the FortiGate GUI.

Field	Description
Type	The type values assigned to FortiGate routes (Static, Connected, RIP, OSPF, or BGP). • **All**: all routes recorded in the routing table • **Connected**:all routes associated with direct connections to FortiGate unit interfaces • **Static**: the static routes that have been added to the routing table manually • **RIP**: all routes learned through RIP. For more information, see RIP on page 120 • **RIPNG**: all routes learned through RIP version 6 (which enables the sharing of routes through IPv6 networks) • **BGP**: all routes learned through BGP. For more information, see BGP on page 199. • **OSPF**: all routes learned through OSPF. For more information, see OSPF on page 158. • **OSPF6**: all routes learned through OSPF version 6 (which enables the sharing of routes through IPv6 networks) • **IS-IS**: all routes learned through IS-IS. For more information, see IS-IS on page 238. • **HA**: RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you are viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster. For more information about HA routing synchronization, see the FortiGate HA User Guide.

Field	Description
Subtype	If applicable, the subtype classification assigned to OSPF routes. An empty string implies an intra-area route. The destination is in an area to which the FortiGate unit is connected. • **OSPF inter area**: the destination is in the OSPF AS, but FortiGate is not connected to that area. • **External 1**: the destination is outside the OSPF AS. This is known as OSPF E1 type. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together. • **External 2**: the destination is outside the OSPF AS. This is known as OSPF E2 type. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. • **OSPF NSSA 1**: same as External 1, but the route was received through a not-so-stubby area (NSSA) • **OSPF NSSA 2**: same as External 2, but the route was received through a not-so-stubby area For more information about OSPF subtypes, see OSPF on page 158.
Network	The IP addresses and network masks of destination networks that FortiGate can reach.
Gateway IP	The IP addresses of gateways to the destination networks.
Interfaces	The interface through which packets are forwarded to the gateway of the destination network.
Up Time	The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable.
Distance	The administrative distance associated with the route. A value of 0 means the route is preferable compared to other routes to the same destination, and the FortiGate unit may routinely use the route to communicate with neighboring routers and access servers. Modifying this distance for dynamic routes is route distribution. See BGP on page 199.

Field	Description
Metric	The metric associated with the route type. The metric of a route influences how the FortiGate unit dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to. **Hop count**: routes learned through RIP **Relative cost**: routes learned through OSPF **Multi-Exit Discriminator (MED)**: routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network. For more information on BGP attributes, see BGP on page 199. By default, the MED value associated with a BGP route is zero. However, the MED value can be modified dynamically. If the value was changed from the default, the Metric column will display a non-zero value. Not displayed when IP version 6 is selected.

Copying DSCP value in GRE tunnels

You can enable an option to allow copying of the DSCP (Differentiated services code point) value in GRE tunnels. This feature enables the keeping of the DSCP marking in the packets after encapsulation for going through GRE tunnels.

To enable DSCP copying, use the following CLI commands:

```
config sys gre-tunnel
   set dscp-copying enable
```

Configuring the maximum number of IP route cache entries

You can configure the maximum number of route cache entries, using the following CLI commands:

```
config system global
   set max-route-cache-size <integer between 0 - 2147483647>
   end
```

Unsetting the field causes the value to be set to the kernel calculated default:

```
config system global
   unset max-route-cache-size
   end
```

Viewing the routing table in the CLI

In the CLI, you can easily view the static routing table just as in the web-based manager or you can view the full routing table.

When viewing the list of static routes using the CLI command `get router static`, it is the configured static routes that are displayed. When viewing the routing table using the CLI command `get router info routing-table all`, it is the entire routing table information that is displayed, including configured and learned routes of all types. The two are different information in different formats.

If VDOMs are enabled on your FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

To view the routing table

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O  - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i  - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
*  - candidate default

S*  0.0.0.0/0 [10/0] via 192.168.183.254, port2
S    1.0.0.0/8 [10/0] via 192.168.183.254, port2
S   2.0.0.0/8 [10/0] via 192.168.183.254, port2
C   10.142.0.0/23 is directly connected, port3
B   10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C 192.168.182.0/23 is directly connected, port2
```

Examining an entry:

```
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

Value	Description
B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route, including netmask.
[20/0]	20 indicates and administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
10.142.0.74	The gateway, or next hop.
port3	The interface used by this route.
2d18h02m	How old this route is. In this case, it is almost three days old.

To view the kernel routing table

```
# get router info kernel

tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0->10.11.201.0/24
    pref=10.11.201.4 gwy=0.0.0.0 dev=5(external1)

tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0->172.20.120.0/24
    pref=172.20.120.146 gwy=0.0.0.0 dev=6(internal)
```

The parts of the routing table entry are:

Value	Description
tab	Table number. This will be either 254 (unicast) or 255 (multicast).
vf	Virtual domain of the firewall. This is the vdom index number. If vdoms are not enabled, this number will be 0.
type	Type of routing connection. Valid values include: 0 - unspecific 1 - unicast 2 - local 3 - broadcast 4 - anycast 5 - multicast 6 - blackhole 7 - unreachable 8 - prohibited
proto	Type of installation. This indicates where the route came from. Valid values include: 0 - unspecific 2 - kernel 11 - ZebOS routing module 14 - FortiOS 15 - HA 16 - authentication based 17 - HA1
prio	Priority of the route. Lower priorities are preferred.
->10.11.201.0/24 **(->x.x.x.x/mask)**	The IP address and subnet mask of the destination
pref	Preferred next hop along this route
gwy	Gateway - the address of the gateway this route will use
dev	Outgoing interface index. This number is associated with the interface for this route, and if VDOMs are enabled the VDOM will be included here as well. If an interface alias is set for this interface, it will also be displayed here.

Searching the routing table

You can apply a filter to search the routing table and display only certain routes. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed. An implicit AND condition is applied to all of the search parameters you specify.

For example, if the FortiGate unit is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select **Connected** from the **Type** list, type `172.16.14.0/24` in the **Network** field, and then select **Apply Filter** to display the associated routing table entry or entries. Any entry that contains the word “Connected” in its **Type** field and the specified value in the **Gateway** field will be displayed.

In this example, you will apply a filter to search for an entry for static route to 10.10.10.10/24.

To search the FortiGate unit routing table - web-based manager

1. Go to **Monitor > Routing Monitor**.
2. From the **Type** list, select the type of route to display. In our example, select **Static**.
3. If you want to display routes to a specific network, type the IP address and netmask of the network in the Networks field. In our example, enter `10.10.10.10/24`.
4. If you want to display routes to a specific gateway, type the IP address of the gateway in the **Gateway** field.
5. Select **Apply Filter**.

All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.

To search the FortiGate unit routing table - CLI

```
FGT # get router info routing-table details 10.10.10.10
Routing entry for 10.10.10.10/24
   Known via "static", distance 10, metric 0, best
```

If there are multiple routes that match your filter, they will all be listed and the best match will be at the top of the list and indicated by the word best.

Building the routing table

In the factory default configuration, the FortiGate unit routing table contains a single static default route. You can add routing information to the routing table by defining additional static routes.

It is possible that the routing table is faced with several different routes to the same destination—the IP addresses of the next-hop router specified in those routes or the FortiGate interfaces associated with those routes may vary. In this situation, the “best” route is selected from the table.

The FortiGate unit selects the “best” route for a packet by evaluating the information in the routing table. The “best” route to a destination is typically associated with the shortest distance between the FortiGate unit and the

closest gateway, also known as a next-hop router. In some cases, the next best route may be selected if the best route is unavailable.

The FortiGate unit installs the best available routes in the unit's forwarding table, which is a subset of the unit's routing table. Packets are forwarded according to the information in the forwarding table.

Static routing security

Securing the information on your company network is a top priority for network administrators. Security is also required as the routing protocols used are internationally known standards that typically provide little or no inherent security by themselves.

The two reasons for securing your network are the sensitive and proprietary information on your network, and also your external bandwidth. Hackers can steal not only your information, but they can also steal your bandwidth. Routing is a good low level way to secure your network, even before UTM features are applied.

Routing provides security to your network in a number of ways including obscuring internal network addresses with NAT and blackhole routing, using RPF to validate traffic sources, and maintaining an access control list (ACL) to limit access to the network.

Network Address Translation

Network address translation (NAT) is a method of changing the address from which traffic appears to originate. This practice is used to hide the IP address on a company's internal networks, and helps prevent malicious attacks that use those specific addresses.

This is accomplished by the router connected to that local network changing all the IP addresses to its externally connected IP address before sending the traffic out to the other networks, such as the Internet. Incoming traffic uses the established sessions to determine which traffic goes to which internal IP address. This also has the benefit of requiring only the router to be very secure against external attacks, instead of the whole internal network, as would be the case without NAT. Securing the network is much cheaper and easier to maintain.

Configuring NAT on your FortiGate unit includes the following steps:

1. Configure your internal network. For example, use the `10.11.101.0` subnet.
2. Connect your internal subnet to an interface on your FortiGate unit. For example, use `port1`.
3. Connect your external connection (for example, an ISP gateway of `172.20.120.2`) to another interface on your Fortigate unit (for example, `port2`).

Configure security policies to allow traffic between port1 and port2 on your FortiGate unit, ensuring that the NAT feature is enabled.

The above steps show that traffic from your internal network will originate on the 10.11.101.0 subnet and pass on to the 172.20.120.0 network. The FortiGate unit moves the traffic to the proper subnet. In doing that, the traffic appears to originate from the FortiGate unit interface on that subnet and it does not appear to originate from where it actually came from.

NAT "hides" the internal network from the external network. This provides security through obscurity. If a hacker tries to directly access your network, they will find the Fortigate unit, but they will not know about your internal network. The hacker would have to get past the security-hardened FortiGate unit to gain access to your internal network. NAT will not prevent hacking attempts that piggy back on valid connections between the internal network and the outside world. However, other UTM security measures can deal with these attempts.

Another security aspect of NAT is that many programs and services have problems with NAT. Consider if someone on the Internet tries to initiate a chat with someone on the internal network. The outsider can access

only the FortiGate unit's external interface, unless the security policy allows the traffic through to the internal network. If allowed in, the correct internal user would respond to the chat. However, if it is not allowed, the request to chat will be refused or it will time out. This is accomplished in the security policy by allowing or denying different protocols.

Access control list

An access control list (ACL) is a table of addresses that have permission to send and receive data over a router's interface or interfaces. The router maintains an ACL, and when traffic comes in on a particular interface it is buffered, while the router checks the ACL to see if that traffic is allowed over that port. If it is allowed on that incoming interface, the next step is to check the ACL for the destination interface. If the traffic also passes that check, the buffered traffic is delivered to its destination. If either of those steps fail the ACL check, the traffic is dropped and an error message may be sent to the sender. The ACL ensures that traffic follows expected paths and any unexpected traffic is not delivered. This stops many network attacks. However, to be effective, the ACL must be kept up to date. When employees or computers are removed from the internal network, their IP addresses must also be removed from the ACL. For more information about the ACL, see the router chapter of the FortiGate CLI Reference.

Blackhole routes

A blackhole route is a route that drops all traffic sent to it. It is very much like /dev/null in Linux programming.

Blackhole routes are used to dispose of packets instead of responding to suspicious inquiries. This provides added security since the originator will not discover any information from the target network.

Blackhole routes can also limit traffic on a subnet. If some subnet addresses are not in use, traffic to those addresses, which may be valid or malicious, can be directed to a blackhole for added security and to reduce traffic on the subnet.

The loopback interface, which is a virtual interface that does not forward traffic, was added to allow easier configuration of blackhole routing. Similar to a regular interface, the loopback interface has fewer parameters to configure and all traffic sent to it stops there. Since it cannot have hardware connection or link status problems, it is always available, making it useful for other dynamic routing roles. Once configured, you can use a loopback interface in security policies, routing, and other places that refer to interfaces. You configure this feature only from the CLI. For more information, see the system chapter of the FortiGate CLI Reference.

Configuring IPv6 blackhole routes

You can configure IPv6 blackhole routes. In the FortiGate GUI, select **Network > Static Routes** and select **Create New**. In the **Interface** field, choose **Blackhole**.

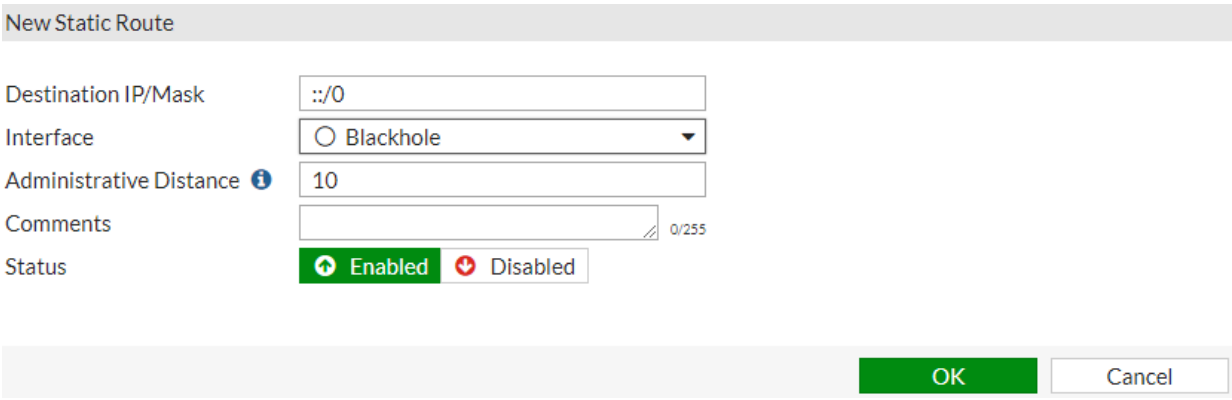

Adding a blackhole route with a priority

You can add a priority to a blackhole route to change its position relative to kernel routes in the routing table.

To add a blackhole route with a priority, use the following CLI commands:

```
config router static
   edit 23
      set blackhole enable
      set priority 200
   end
```

IPv6 blackhole static routing

System administrators use blackhole routing to divert unwanted traffic, such as packets from a Denial of Service (DoS) attack or communications from an illegal source. The traffic is routed to a dead interface, or a host designed to collect information for investigation. This mitigates the impact of the attack on the network.

You can enable the use of blackhole routing, using the following CLI commands:

```
config router static6
   edit <ID #>
      set blackhole enable
   end
```

Reverse path lookup

Whenever a packet arrives at one of the FortiGate unit's interfaces, the unit determines whether the packet was received on a legitimate interface by doing a reverse lookup using the source IP address in the packet header. This is also called anti-spoofing. If the FortiGate unit cannot communicate with the computer at the source IP address through the interface on which the packet was received, the FortiGate unit drops the packet as it is likely a hacking attempt.

If the destination address can be matched to a local address, and the local configuration permits delivery, the FortiGate unit delivers the packet to the local network. If the packet is destined for another network, the Fortigate unit forwards the packet to a next-hop router according to a policy route and the information stored in the FortiGate forwarding table.

Multipath routing and determining the best route

Multipath routing occurs when more than one entry to the same destination is present in the routing table. When multipath routing happens, the FortiGate unit may have several possible destinations for an incoming packet, forcing the FortiGate unit to decide which next-hop is the best one.

It should be noted that some IP addresses will be rejected by routing protocols. These are called Martian addresses. They are typically IP addresses that are invalid and not routable because they have been assigned an address by a misconfigured system, or are spoofed addresses.

Two methods to manually resolve multiple routes to the same destination are to lower the administrative distance of one route or to set the priority of both routes. For the FortiGate unit to select a primary (preferred) route, manually lower the administrative distance associated with one of the possible routes. Setting the priority on the routes is a FortiGate unit feature and may not be supported by non-Fortinet routers.

Administrative distance is based on the expected reliability of a given route. It is determined through a combination of the number of hops from the source and the protocol used. A hop is when traffic moves from one

router to the next. More hops from the source means more possible points of failure. The administrative distance can be from 1 to 255, with lower numbers being preferred. A distance of 255 is seen as infinite and will not be installed in the routing table.

Here is an example to illustrate how administration distance works. If there are two possible routes traffic can take between two destinations with administration distances of 5 (always up) and 31 (sometimes not available), the traffic will use the route with an administrative distance of 5. If for some reason the preferred route (admin distance of 5) is not available, the other route will be used as a backup.

Different routing protocols have different default administrative distances. These different administrative distances are based on a number of factors of each protocol such as reliability, speed, and so on. The default administrative distances for any of these routing protocols are configurable.

Default administrative distances for routing protocols and connections

Routing protocol	Default administrative distance
Direct physical connection	1
Static	10
EBGP	20
OSPF	110
IS-IS	115
RIP	120
IBGP	200

Another method to determine the best route is to manually change the priority of both routes in question. If the next-hop administrative distances of two routes on the FortiGate unit are equal, it may not be clear which route the packet will take. Manually configuring the priority for each of those routes will make it clear which next-hop will be used in the case of a tie. The priority for a route can be set in the CLI, or when editing a specific static route, as described in the next section. Lower priority routes are preferred. Priority is a Fortinet value that may or may not be present in other brands of routers.

All entries in the routing table are associated with an administrative distance. If the routing table contains several entries that point to the same destination (the entries may have different gateways or interface associations), the FortiGate unit compares the administrative distances of those entries first, selects the entries having the lowest distances, and installs them as routes in the FortiGate unit forwarding table. As a result, the FortiGate unit forwarding table contains only those routes that have the lowest distances to every possible destination. While only static routing uses administrative distance as its routing metric, other routing protocols, such as RIP, can use metrics that are similar to administrative distance.

Route priority

After the FortiGate unit selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can configure the priority field through the CLI or the web-based manager. Priority values can range from 0 to 4 294 967 295. The route with the lowest value in the priority field is considered the best route. It is also the primary route.

To change the priority of a route - web-based manager

1. Go to **Network > Static Routes**.
2. Select the route entry, and select **Edit**.
3. Select **Advanced Options**.
4. Enter the **Priority** value.
5. Select **OK**.

To change the priority of a route - CLI

The following command changes the priority to 5 for a route to the address `10.10.10.1` on the `port1` interface.

```
config router static
   edit 1
      set device port1
      set gateway 10.10.10.10
      set dst 10.10.10.1
      set priority 5
   end
```

If there are other routes set to priority 10, the route set to priority 5 will be preferred. If there are routes set to priorities less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal-cost multi-path (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, different sessions will resolve this problem by using different routes to the same address.

Use of firewall addresses for static route destinations

To help prevent false positive when scanning for duplicate static routes, the dst_addr field is also checked.

Removing RPF checks from the state evaluation process

You can remove RPF (reverse path forwarding) state checks without needing to enable asymmetric routing. You can disable state checks for traffic received on specific interfaces.

Disabling state checks makes a FortiGate unit less secure and should only be done with caution.

To remove RPF checks from the state evaluation process, use the following CLI commands:

```
config system interface
   edit <interface_name>
      set src-check disable
   end
```

Troubleshooting static routing

When there are problems with your network that you believe to be related to static routing, there are a few basic tools available to locate the problem.

These tools include:

- Ping
- Traceroute
- Examine routing table contents

Ping

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is no packet loss detected, your basic network connectivity is OK.

If there is some packet loss detected, you should investigate:

- Possible ECMP, split horizon, network loops
- Cabling to ensure no loose connections

If there is total packet loss, you should investigate:

- Hardware: ensure cabling is correct, and all equipment between the two locations is accounted for
- Addresses and routes: ensure all IP addresses and routing information along the route is configured as expected
- Firewalls: ensure all firewalls are set to allow PING to pass through

To ping from a Windows PC

1. Go to a DOS prompt. Typically you go to **Start > Run**, enter `cmd` and select **OK**.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate unit with four packets.

To ping from an Apple computer

1. Open the Terminal.
2. Enter `ping 10.11.101.100`.
3. If the ping fails, it will stop after a set number of attempts. If it succeeds, it will continue to ping repeatedly. Press `Control+C` to end the attempt and see gathered data.

To ping from a Linux PC

1. Go to a command line prompt.
2. Enter "`/bin/etc/ping 10.11.101.101`".

Traceroute

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, traceroute can be used to locate exactly where the problem is.

To use traceroute on a Windows PC

1. Go to a DOS prompt. Typically you go to **Start > Run**, enter "`cmd`" and select **OK**.
2. Enter "`tracert fortinet.com`" to trace the route from the PC to the Fortinet website.

To use traceroute from an Apple computer

1. Open the Terminal.
2. Enter `traceroute fortinet.com`.
3. The terminal will list the number of steps made. Upon reaching the destination, it will list three asterisks per line. Press `Control+C` to end the attempt.

To use traceroute on a Linux PC

1. Go to a command line prompt.
2. Enter "`/bin/etc/traceroute fortinet.com`".
   The Linux traceroute output is very similar to the MS Windows traceroute output.

Examine routing table contents

The first place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route is not used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. Note that if your FortiGate unit is in Transparent mode, you will not be able to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the web-based manager, use the Routing Monitor. Go to **Monitor > Routing Monitor**. In the CLI, use the command `get router info routing-table all`.

Static routing tips

When your network goes beyond basic static routing, here are some tips to help you plan and manage your static routing.

Always configure a default route

The first thing you configure on a router on your network should be the default route. And where possible the default routes should point to either one or very few gateways. This makes it easier to locate and correct problems in the network. By comparison, if one router uses a second router as its gateway which uses a fourth for

its gateway and so on, one failure in that chain will appear as an outage for all the devices downstream. By using one or very few addresses as gateways, if there is an outage on the network it will either be very localized or network-wide. Either outage is easy to troubleshoot.

Have an updated network plan

A network plan lists different subnets, user groups, and different servers. Essentially, it puts all your resources on the network and shows how the parts of your network are connected. Keeping your plan updated will also help you troubleshoot problems more quickly when they arise.

A network plan helps your static routing by eliminating potential bottlenecks and helping troubleshoot any routing problems that come up. Also, you can use it to plan for the future and act on any changes to your needs or resources more quickly.

Plan for expansion

No network remains the same size. At some time, all networks grow. If you take future growth into account, there will be less disruption to your existing network when that growth happens. For example, allocating a block of addresses for servers can easily prevent having to re-assign IP addresses to multiple servers due to a new server.

With static routing, if you group parts of your network properly you can easily use network masks to address each part of your network separately. This will reduce the amount of administration required both to maintain the routing and to troubleshoot any problems.

Configure as much security as possible

Securing your network through static routing methods is a good low level method to defend both your important information and your network bandwidth.

- Implement NAT to obscure your IP address is an excellent first step
- Implement black hole routing to hide which IP addresses are in use or not on your local network
- Configure and use access control list (ACL) to help ensure you know only valid users are using the network

All three features limit access to the people who should be using your network and obscure your network information from the outside world and potential hackers.

Policy routing

Policy routing enables you to redirect traffic away from a static route. This can be useful if you want to route certain types of network traffic differently. You can use incoming traffic's protocol, source address or interface, destination address, or port number to determine where to send the traffic. For example, generally network traffic would go to the router of a subnet, but you might want to direct SMTP or POP3 traffic directly to the mail server on that subnet.

If you have configured the FortiGate unit with routing policies and a packet arrives at the FortiGate unit, the FortiGate unit starts at the top of the Policy Route list and attempts to match the packet with a policy. If a match is found and the policy contains enough information to route the packet (a minimum of the IP address of the next-hop router and the FortiGate interface for forwarding packets to it), the FortiGate unit routes the packet using the information in the policy. If no policy route matches the packet, the FortiGate unit routes the packet using the routing table.

Most policy settings are optional, and a matching policy alone might not provide enough information for forwarding the packet. In fact, the FortiGate almost always requires a matching route in the routing table in order to use a policy route. The FortiGate unit will refer to the routing table in an attempt to match the information in the packet header with a route in the routing table.

Policy route options define which attributes of a incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the FortiGate unit routes the packet through the specified interface to the specified gateway.

To view policy routes go to **Network > Policy Routes**.

Field	Description
Create New	Add a policy route. See Adding a policy route on page 90.
Edit	Edit the selected policy route.
Delete	Delete the selected policy route.
Move To	Move the selected policy route. Enter the new position and select **OK**. For more information, see Moving a policy route on page 93.
#	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.
Incoming	The interfaces on which packets subjected to route policies are received.
Outgoing	The interfaces through which policy routed packets are routed.
Source	The IP source addresses and network masks that cause policy routing to occur.
Destination	The IP destination addresses and network masks that cause policy routing to occur.

Adding a policy route

To add a policy route, go to **Network > Policy Routes** and select **Create New**.

Field	Description
Protocol	Select from existing or specify the protocol number to match. The Internet Protocol Number is found in the IP packet header. RFC 5237 describes protocol numbers and you can find a list of the assigned protocol numbers here. The range is from 0 to 255. A value of `0` disables the feature. Commonly used **Protocol** settings include 6 for TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.
Incoming Interface	Select the name of the interface through which incoming packets subjected to the policy are received.
Source Address / Mask	To perform policy routing based on IP source address, type the source address and network mask to match. A value of `0.0.0.0/0.0.0.0` disables the feature.
Destination Address / Mask	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of `0.0.0.0/0.0.0.0` disables the feature.
Destination Ports	To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature. The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.
Type of Service	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see Type of Service on page 92.
Outgoing Interface	Select the name of the interface through which packets affected by the policy will be routed.
Gateway Address	Type the IP address of the next-hop router that the FortiGate unit can access through the specified interface.

Example policy route

Configure the following policy route to send all FTP traffic received at `port1` out the `port10` interface and to a next hop router at IP address `172.20.120.23`. To route FTP traffic, set protocol to 6 (for TCP) and set both of the destination ports to 21 (the FTP port).

Field	Value
Protocol	6

Field	Value
Incoming interface	port1
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0
Destination Ports	From 21 to 21
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)
Outgoing interface	port10
Gateway Address	172.20.120.23

Enabling or disabling individual policy routes

You can enable or disable individual policy routes.

To configure IPv4 policy routes, use the following CLI commands:

```
config router policy
   edit 0
      set status [enable|disable]
      end
```

To configure IPv6 policy routes, use the following CLI commands:

```
config router policy6
   edit 0
      set status [enable|disable]
      end
```

Type of Service

Type of service (TOS) is an 8-bit field in the IP header that allows you to determine how the IP datagram should be delivered, with such qualities as delay, priority, reliability, and minimum cost.

Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority TOS is 0, the highest is 7 - when bits 3, 4, and 5 are all set to 1. The router tries to match the TOS of the datagram to the TOS on one of the possible routes to the destination. If there is no match, the datagram is sent over a zero TOS route.

Using increased quality may increase the cost of delivery because better performance may consume limited network resources. For more information, see RFC 791 and RFC 1349.

The role of each bit in the IP header TOS 8-bit field

Bit	Quality	Description
bits 0, 1, 2	**Precedence**	Some networks treat high precedence traffic as more important traffic. Precedence should only be used within a network, and can be used differently in each network. Typically you do not care about these bits.
bit 3	**Delay**	When set to 1, this bit indicates low delay is a priority. This is useful for such services as VoIP where delays degrade the quality of the sound.
bit 4	**Throughput**	When set to 1, this bit indicates high throughput is a priority. This is useful for services that require lots of bandwidth, such as video conferencing.
bit 5	**Reliability**	When set to 1, this bit indicates high reliability is a priority. This is useful when a service must always be available, such as with DNS servers.
bit 6	**Cost**	When set to 1, this bit indicates low cost is a priority. Generally there is a higher delivery cost associated with enabling bits 3, 4, or 5, and bit 6 indicates to use the lowest cost route.
bit 7	**Reserved for future use**	Not used at this time.

For example, if you want to assign low delay and high reliability for a VoIP application, where delays are unacceptable, you would use a bit pattern of xxx1x1xx where ‘x’ indicates that bit can be any value. Since all bits are not set, this is a good use for the bit mask. If the mask is set to 0x14, it will match any TOS packets that are set to low delay and high reliability.

Moving a policy route

A routing policy is added to the bottom of the routing table when it is created. If you prefer to use one policy over another, you may want to move it to a different location in the routing policy table.

The option to use one of two routes happens when both routes are a match, for example `172.20.0.0/255.255.0.0` and `172.20.120.0/255.255.255.0`. If both of these routes are in the policy table, both can match a route to `172.20.120.112` but you would consider the second one a better match. In that case, the best match route should be positioned before the other route in the policy table.

To change the position of a policy route in the table, go to **Network > Policy Routes** and select **Move To** for the policy route you want to move.

Field	Description
Before/After	Select Before to place the selected policy route before the indicated route. Select After to place it following the indicated route.
Policy route ID	Enter the Policy route ID of the route in the Policy route table to move the selected route before or after.

Use of firewall addresses for policy route destinations

When you configure a policy route, you can use firewall addresses and address groups. The only exception for address types that can be used is the URL type of address object.

Transparent mode static routing

FortiOS operating modes allow you to change the configuration of your FortiGate unit depending on the role it needs to fill in your network.

NAT/Route operating mode is the standard mode, where all interfaces are accessed individually and traffic can be routed between ports to travel from one network to another.

In transparent operating mode, all physical interfaces act like one interface. The FortiGate unit essentially becomes a bridge. Traffic coming in over any interface is broadcast back out over all the interfaces on the FortiGate unit.

In transparent mode, there is no entry for routing at the main level of the menu on the web-based manager display as there is in NAT/Route mode. Routing is instead accessed through the network menu option.

To view the routing table in transparent mode, go to **Network > Routing Table**.

When viewing or creating a static route entry in transparent mode, there are only three fields available.

Field	Description
Destination IP / Mask	The destination of the traffic being routed. The first entry is attempted first for a match, then the next, and so on until a match is found or the last entry is reached. If no match is found, the traffic will not be routed. Use 0.0.0.0 to match all traffic destinations. This is the default route.
Gateway	Specifies the next hop for the traffic. Generally the gateway is the address of a router on the edge of your network.
Priority	The priority is used if there is more than one match for a route. This allows multiple routes to be used, with one preferred. If the preferred route is unavailable, the other routes can be used instead. Valid range of priority can be from 0 to 4 294 967 295. If more than one route matches and they have the same priority, it becomes an ECMP situation and traffic is shared among those routes. See Transparent mode static routing on page 94.

When configuring routing on a FortiGate unit in transparent mode, remember that all interfaces must be connected to the same subnet. That means all traffic will be coming from and leaving on the same subnet. This is important because it limits your static routing options to only the gateways attached to this subnet. For example, if you only have one router connecting your network to the Internet, all static routing on the FortiGate unit will use that gateway. For this reason, static routing on FortiGate units in transparent mode may be a bit different, but it is not as complex as routing in NAT/Route mode.

Static routing example

This is an example of a typical small network configuration that uses only static routing.

This network is in a dental office that includes a number of dentists, assistants, and office staff. The size of the office is not expected to grow significantly in the near future, and the network usage is very stable (there are no new applications being added to the network).

The users on the network are:

- Administrative staff: access to local patient records, and perform online billing
- Dentists: access and update local patient records, research online from desk
- Assistants: access and update local patient records in exam rooms

The distinction here is mainly that only the administrative staff and dental office need access to the Internet. All the other traffic is local and does not need to leave the local network. Routing is only required for the outbound traffic, and the computers that have valid outbound traffic.

Configuring routing only on computers that need it, acts as an additional layer of security by helping prevent malicious traffic from leaving the network.

Network layout and assumptions

The computers on the network are administrative staff computers, dental office computers, and dental exam room computers. While there are other devices on the local network, such as printers, they do not need Internet access or any routing.

This networked office equipment includes 1 administrative staff PC, 3 dentist's PCs, and 5 exam room PCs. There is also a network printer and a router on the network.

Assumptions about these computers and network include:

- The FortiGate unit is a model with interfaces labeled port1 and port2.
- The FortiGate unit has been installed and is configured in NAT/Route mode.
- VDOMs are not enabled.
- The computers on the network are running MS Windows software.
- Any hubs required in the network are not shown in the network diagram.
- The network administrator has access to the ISP IP addresses and is the super_admin administrator on the FortiGate unit.

Static routing example device names, IP addresses, and level of access

Device name	IP address	Need external access?
Router	192.168.10.1	YES
Admin	192.168.10.11	YES

Device name	IP address	Need external access?
Dentist1-3	192.168.10.21-23	YES
Exam1-5	192.168.10.31-35	NO
Printer	192.168.10.41	NO

General configuration steps

The steps to configuring routing on this network are:

1. Get your ISP information such as DNS, gateway, etc.
2. Configure FortiGate unit
3. Configure administrator PC and dentists' PCs
4. Testing network configuration

Get your ISP information such as DNS, gateway, etc.

Your local network connects to the Internet through your Internet Service Provider (ISP). They have IP addresses that you need to configure your network and routing.

The addresses needed for routing are your assigned IP address, DNS servers, and the gateway.

Configure FortiGate unit

The FortiGate unit will have two interfaces in use: one connected to the internal network, and one connected to the external network. Port1 will be the internal interface and port2 will be the external interface.

To configure the FortiGate unit:

1. Configure the internal interface (port1)
2. Configure the external interface (port2)
3. Configure networking information
4. Configure basic security policies
5. Configure static routing

Configure the internal interface (port1)

To configure the internal interface (port1) - web based manager

1. Go to **Network > Interfaces**. Highlight **port1** and select **Edit**.
2. Enter the following:

Addressing Mode	Manual
IP/Netmask	172.100.1.1/255.255.255.0

Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port1) - CLI

```
config system interface
   edit port1
      set IP 192.168.10.1 255.255.255.0
      set allowaccess https ping telnet
      set description "internal network"
   end
end
```

Configure the external interface (port2)

The external interface connects to your ISP's network. You need to know the IP addresses in their network that you should connect to. In this example, the address that the ISP gave you is 172.100.20.20, which will connect to the gateway at 172.100.20.5 on their network, and their DNS servers are 172.11.22.33 and 172.11.22.34.

To configure the internal interface (port2) - web based manager

1. Go to **Network > Interfaces**. Highlight **port2** and select **Edit**.
2. Enter the following:

Addressing Mode	Manual
IP/Netmask	172.100.20.20/255.255.255.0
Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port2) - CLI

```
configure system interface
   edit port2
      set IP 172.100.20.20 255.255.255.0
      set allowaccess https ping telnet
      set description "internal network"
   end
end
```

Configure networking information

Networking information includes the gateway and DNS servers. Your FortiGate unit requires a connection to the Internet for antivirus and other periodic updates.

To configure networking information - web-based manager

1. Go to **Network > DNS**.
2. Enter the primary and secondary DNS addresses.

3. Select **Apply**.

To configure networking information - CLI

```
config system global
   set dns_1 172.11.22.33
   set dns_2 172.11.22.34
end
```

Configure basic security policies

For traffic to flow between the internal and external ports in both directions, as a minimum, two security policies are required. More can be used to further limit or direct traffic, as needed, but will not be included here.

Before configuring the security policies, a firewall address group is configured for the PCs that are allowed Internet access. This prevents a PC without Internet privileges from accessing the Internet.

The security policy assumptions are:

- Only the basic networking services have been listed as allowed, for added security. Others can easily be added as the users require them.
- In this example, to keep things simple, both incoming and outgoing security policies are the same. In a real network there are applications that are allowed out but not in, and vice versa.
- Endpoint control has been enabled to ensure that all computers on the local network are running FortiClient and those installs are up to date. This feature ensures added security on your local network without the need for the network administrator to continually bother users to update their software. The FortiGate unit can store an up to date copy of the FortiClient software and offer a URL to it for users to install it if they need to.

To configure security policies - web-based manager

1. Go to **Policy & Objects > Objects > Addresses**.
2. Create a new Firewall Address entry for each of:

PC Name	IP Address	Interface
Admin	192.168.10.11	port1
Dentist1	192.168.10.21	port1
Dentist2	192.168.10.22	port1
Dentist3	192.168.10.23	port1

3. Go to **Policy & Objects > Objects > Addresses**.
4. Select the dropdown arrow next to **Create New** and select **Address Group**.
5. Name the group Internet_PCs.
6. Add Admin, Dentist1, Dentist2, and Dentist3 as members of the group.
7. Select **OK**.
8. Go to **Policy & Objects > Policy > IPv4**.
9. Select **Create New**.
10. Enter the following: DH - port2(external) -> port1(internal)

Incoming Interface	port2
Source Address	all
Outgoing Interface	port1
Destination Address	Internet_PCs
Schedule	always
Service	Multiple. Select **DHCP, DNS,FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.**
Action	ACCEPT
Log Allowed Traffic	Enabled

11. Select **OK**.
12. Select **Create New**.
13. Enter the following:

Incoming Interface	port1
Source Address	Internet_PCs
Outgoing Interface	port2
Destination Address	all
Schedule	always
Service	Multiple. Select **DHCP, DNS,FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.**
Action	ACCEPT
Log Allowed Traffic	Enabled

14. Select **OK**.

To configure security policies - CLI

```
config firewall address
   edit "Admin"
      set associated-interface "port1"
      set subnet 192.168.10.11 255.255.255.255
   next
   edit "Dentist1"
      set associated-interface "port1"
```

```
        set subnet 192.168.10.21 255.255.255.255
      next
      edit "Dentist2"
        set associated-interface "port1"
        set subnet 192.168.10.22 255.255.255.255
      next
      edit "Dentist3"
        set associated-interface "port1"
        set subnet 192.168.10.23 255.255.255.255
      end
    config firewall addrgrp
      edit Internet_PCs
        set member Admin Dentist1 Dentist2 Dentist3
      end
    config firewall policy
      edit 1
        set srcintf port1
        set dstintf port2
        set srcaddr Internet_PCs
        set dstaddr all
        set action accept
        set schedule always
        set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3" "SMTP" "SSH"
        set logtraffic enable
        set label "Section2"
        set endpoint-restrict-check no-av db-outdated
      next
      edit 2
        set srcintf port2
        set dstintf port1
        set srcaddr all
        set dstaddr Internet_PCs
        set action accept
        set schedule always
        set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3" "SMTP" "SSH"
        set logtraffic enable
        set label "Section2"
        set endpoint-restrict-check no-av db-outdated
      end
    end
```

Adding FortiClient enforcement to interfaces

You can enforce the use of FortiClient on individual interfaces.

In the FortiGate GUI, select **Network** > **Interfaces** and choose an interface. Under the **Admission Control** heading, you can enable the **Allow FortiClient Connections** setting. Once you enable this setting, two more options become visible: **Discover Clients (Broadcast)** and **FortiClient Enforcement**. When you enable FortiClient enforcement, you enforce that in order for incoming traffic to pass through that interface, it must be initiated by a device running FortiClient.

Once you enforce the use of FortiClient on the interface, you should also configure FortiClient profiles for the incoming connections. You can also set up any exemptions that are needed. Just below the **FortiClient Enforcement** option are fields for **Exempt Sources** and **Exempt Destinations/Services**. These can be selected from address or services objects already configured on the FortiGate.

In the CLI, use the following commands:

```
config system interface
   edit port1
      set listen-forticlient-connection [enable|disable]
      set endpoint-compliance [enable|disable]
      end
```

Configure static routing

With the rest of the FortiGate unit configured, static routing is the last step before moving on to the rest of the local network. All traffic on the local network will be routed according to this static routing entry.

To configure Fortinet unit static routing - web-based manager

1. Go to **Network > Static Routes**.
2. Select the top route on the page and then select **Edit**.
3. Enter the following information:

Destination IP/Mask	172.100.20.5
Gateway	172.100.20.5
Interface	port2
Administrative Distance	10

4. Select **OK**.

To configure Fortinet unit static routing - CLI

```
configure routing static
   edit 1
      set gateway 172.100.20.5
      set distance 10
      set device port2
      set dst 0.0.0.0
   end
end
```

Configure administrator PC and dentists' PCs

After the router is configured, we need to configure the computers that require Internet access. These computers need routing to be configured on them. As the other computers do not require routing, they are not included here.

The procedure to configure these computers is the same. Repeat the following procedure for the corresponding PCs.

The Windows CLI procedure does not configure the DNS entries. It just adds the static routes.

To configure routing and DNS on administrator and dentists' PCs - Windows GUI

1. On the PC, select **Start > Control Panel > Network Connections**.
2. Right click on the network connection to your local network that has a status of Connected, and select **Properties**.
3. Under the **General** tab, from the list select **TCP/IP**, and **Properties**.
4. Under **Gateway**, enter the FortiGate unit address (192.168.10.1).
5. Enter the primary and secondary DNS server addresses from your ISP (172.11.22.33 and 172.11.22.34).
6. Select **OK**.

To configure routing on administrator and dentists' PCs - Windows CLI

1. On the PC, select **Start > Run**, enter "`cmd`", and select **OK**.
2. At the command prompt, type:
```
route ADD 0.0.0.0 MASK 0.0.0.0 172.100.20.5 METRIC 10
route ADD 192.168.10.0 MASK 255.255.255.0 192.168.10.1 METRIC 5
```
3. Confirm these routes have been added. Type:
```
route PRINT
```
   If you do not see the two routes you added, try adding them again, while paying attention to avoid spelling mistakes.
4. Test that you can communicate with other computers on the local network, and with the Internet. If there are no other computers on the local network, connect to the FortiGate unit.

Configure other PCs on the local network

The PCs on the local network without Internet access (for example, the exam room PCs) can be configured now.

As this step does not require any routing, details have not been included.

Testing network configuration

There are three tests to run on the network to ensure proper connectivity:

- To test that PCs on the local network can communicate
- Test that Internet_PCs on the local network can access the Internet
- Test that non-Internet_PCs cannot access the Internet

Test that PCs on the local network can communicate

1. Select any two PCs on the local network, such as Exam4 and Dentist3.
2. On the Exam4 PC, at the command prompt, enter `ping 192.168.10.23.`
   The output from this command should appear similar to the following:
```
Pinging 192.168.10.23 with 32 bytes of data:

Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
```
3. At the command prompt, enter `exit` to close the window.
4. On the Dentist3 PC, at the command prompt, enter `ping 192.168.10.34.`

The output from this command should appear similar to the following:

```
Pinging 192.168.10.34 with 32 bytes of data:

Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
```

5. At the command prompt, enter `exit` to close the window.
6. Repeat these steps for all PCs on the local network.

   If the output does not appear similar to above, there is a problem with the network configuration between these two PCs.

To test that Internet_PCs on the local network can access the Internet

The easiest way to access the Internet is with an Internet browser. However, if that does not work, it is best to do a traceroute to see at what point the problem is. This can help determine if it is a networking problem such as cabling, or if it is an access problem, such as this PC not having Internet access.

1. Select any PC on the local network that is supposed to have Internet access, such as Admin.
2. On the Admin PC, open an Internet browser and attempt to access a website on the Internet, such as http://www.fortinet.com.

   If this is successful, this PC has Internet access.
3. If step2 was not successful, at the command prompt on the PC, enter `traceroute 22.11.22.33`.

   The output from this command should appear similar to:

```
Pinging 22.11.22.33 with 32 bytes of data:

Reply from 22.11.22.33: bytes=32 time<1m TTL=255
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

Dynamic routing overview

This section provides an overview of dynamic routing and how it compares to static routing.

What is dynamic routing?

Dynamic routing uses a dynamic routing protocol to automatically select the best route to put into the routing table. Instead of having to manually enter static routes in the routing table, dynamic routing automatically receives routing updates and dynamically decides which routes are best to go into the routing table. It is this intelligent and hands-off approach that makes dynamic routing so useful.

Dynamic routing protocols vary in many ways and this is reflected in the various administrative distances assigned to routes learned from dynamic routing. These variations take into account differences in reliability, speed of convergence, and other similar factors. For more information about these administrative distances, see Advanced static routing on page 72.

Comparing static and dynamic routing

A common term used to describe dynamic routing is convergence. Convergence is the ability to work around network problems and outages, for the routing to come together despite obstacles. For example, if the main router between two end points goes down, convergence is the ability to find a way around that failed router and reach the destination. Static routing has zero convergence beyond trying the next route in its limited local routing table. If a network administrator does not fix a routing problem manually, it may never be fixed and may result in a downed network. Dynamic routing solves this problem by involving routers along the route in the decision-making process about the optimal route, and using the routing tables of these routers to find potential routes around the outage. In general, dynamic routing has better scalability, robustness, and convergence. However, the cost of these added benefits includes more complexity and some overhead. For example, the routing protocol uses some bandwidth for its own administration.

Comparing static and dynamic routing

Feature	Static routing	Dynamic routing
Hardware support	Supported by all routing hardware	May require special, more expensive routers
Router memory required	Minimal	Can require considerable memory for larger tables
Complexity	Simple	Complex
Overhead	None	Varying amounts of bandwidth used for routing protocol updates
Scalability	Limited to small networks	Very scalable, better for larger networks

Feature	Static routing	Dynamic routing
Robustness	None: if a route fails, it has to be fixed manually	Robust: traffic routed around failures automatically
Convergence	None	Varies from good to excellent

Dynamic routing protocols

A dynamic routing protocol is an agreed-on method of routing that the sender, receiver, and all routers along the path (route), support. Typically, the routing protocol involves a process running on all computers and routers along that route to enable each router to handle routes in the same way as the others. The routing protocol determines how the routing tables are populated along that route, how the data is formatted for transmission, and what information about a route is included with that route. For example, RIP and BGP use distance vector algorithms and OSPF uses a shortest path first algorithm. Each routing protocol has different strengths and weaknesses. One protocol may have fast convergence, while another may be very reliable, and a third may be very popular for certain businesses like Internet Service Providers (ISPs).

Dynamic routing protocols are different from each other in a number of ways, such as:

- Classful versus classless routing protocols
- Interior versus exterior routing protocols
- Distance vector versus link-state protocols

Classful versus classless routing protocols

Classful and classless routing refers to how the routing protocol handles the IP addresses. In classful addresses, there is the specific address and the host address of the server that address is connected to. Classless addresses use a combination of IP address and netmask.

Classless Inter-Domain Routing (CIDR) was introduced in 1993 (originally with RFC 1519 and most recently with RFC 4632) to keep routing tables from getting too large. With classful routing, each IP address requires its own entry in the routing table. With classless routing, a series of addresses can be combined into one entry, potentially saving vast amounts of space in routing tables.

Current routing protocols that support classless routing, out of necessity, include RIPv2, BGP, IS-IS, and OSPF. Older protocols, such as RIPv1, do not support CIDR addresses.

Interior versus exterior routing protocols

The names interior and exterior and are very descriptive. Interior routing protocols are designed for use within a contained network of limited size, whereas exterior routing protocols are designed to link multiple networks together. They can be used in combination in order to simplify network administration. For example, a network can be built with only border routers of a network running the exterior routing protocol, while all the routers on the network run the interior protocol. This prevents them from connecting outside the network without passing through the border. Exterior routers in such a configuration must have both exterior and interior protocols to communicate with the interior routers and outside the network.

Nearly all routing protocols are interior routing protocols. Only BGP is commonly used as an exterior routing protocol.

You may see interior gateway protocol (IGP) used to refer to interior routing protocols, and exterior gateway protocol (EGP) used to refer to interior routing protocols.

Distance vector versus link-state protocols

Every routing protocol determines the best route between two addresses using a different method. However, there are two main algorithms for determining the best route: distance vector and link-state.

Distance vector protocols

In distance vector protocols, routers are told about remote networks through neighboring routers. The distance part refers to the number of hops to the destination and, in more advanced routing protocols, these hops can be weighted by factors such as available bandwidth and delay. The vector part determines which router is the next step along the path for this route. This information is passed along from neighboring routers with routing update packets that keep the routing tables up to date. Using this method, an outage along a route is reported back along to the start of that route, ideally before the outage is encountered.

On distance vector protocols, RFC 1058, which defines RIP v1, states the following:

> *Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.*

There are four main weaknesses inherent in the distance vector method. Firstly, the routing information is not discovered by the router itself, but is instead reported information that must be relied on to be accurate and up-to-date. The second weakness is that it can take a while for the information to make its way to all the routers who need the information; in other words, it can have slow convergence. The third weakness is the amount of overhead involved in passing these updates all the time. The number of updates between routers in a larger network can significantly reduce the available bandwidth. The fourth weakness is that distance vector protocols can end up with routing-loops. Routing loops are when packets are routed forever around a network, and often occur with slow convergence. The bandwidth required by these infinite loops will slow your network to a halt. There are methods of preventing these loops, however, so this weakness is not as serious as it may first appear.

Link-state protocols

Link-state protocols are also known as shortest path first protocols. Where distance vector uses information passed along that may or may not be current and accurate, in link-state protocols each router passes along information only about the networks and devices that are directly connected to it. This results in a more accurate picture of the network topology around your router, allowing it to make better routing decisions. This information is passed between routers using link-state advertisements (LSAs). To reduce the overhead, LSAs are only sent out when information changes, compared to distance vector sending updates at regular intervals even if no information has changed. The more accurate network picture in link-state protocols greatly speed up convergence and avoid problems such as routing-loops.

Minimum configuration for dynamic routing

Dynamic routing protocols do not pay attention to routing updates from other sources, unless you specifically configure them to do so using CLI redistribute commands within each routing protocol.

The minimum configuration for any dynamic routing to function is to have dynamic routing configured on one interface on the FortiGate unit, and one other router configured as well. Some protocols require larger networks to function as designed.

Minimum configuration based on dynamic protocol

	BGP	RIP	OSPF / IS-IS
Interface	Yes	Yes	Yes
Network	Yes	Yes	Yes
AS	Local and neighbor	No	Yes
Neighbors	At least one	At least one	At least one
Version	No	Yes	No
Router ID	No	No	Yes

Comparison of dynamic routing protocols

Each dynamic routing protocol was designed to meet a specific routing need. Each protocol does some things well, and other things not so well. For this reason, choosing the right dynamic routing protocol for your situation is not an easy task.

Features of dynamic routing protocols

Each protocol is better suited for some situations over others.

Choosing the best dynamic routing protocol depends on the size of your network, speed of convergence required, the level of network maintenance resources available, what protocols the networks you connect to are using, and so on. For more information about these dynamic routing protocols, see RIP on page 120, BGP on page 199, OSPF on page 158, and IS-IS on page 238.

Comparing RIP, BGP, and OSPF dynamic routing protocols

Protocol	RIP	BGP	OSPF / IS-IS
Routing algorithm	Distance vector, basic	Distance vector, advanced	Link-state
Common uses	Small, non-complex networks	Network backbone, ties multinational offices together	Common in large, complex enterprise networks

Protocol	RIP	BGP	OSPF / IS-IS
Strengths	Fast and simple to implement Near universal support Good when no redundant paths	Graceful restart BFD support Only needed on border routers Summarize routes	Fast convergence Robust Little management overhead No hop count limitation Scalable
Weakness	Frequent updates can flood network Slow convergence Maximum 15 hops may limit network configuration	Required full mesh in large networks can cause floods Route flap Load-balance multi-homed networks Not available on low-end routers	Complex No support for unequal cost multipath routing Route summary can require network changes
Authentication	Optional authentication using text string or MD5 password. (RIP v1 has no authentication)		
IPv6 support	Only in RIPng	Only in BGP4+	Only in OSPF6 / Integrated IS-IS

Routing protocols

- **Routing Information Protocol (RIP)** uses classful routing, as well as incorporating various methods to stop incorrect route information from propagating, such as the poisoned horizon method. However, on larger networks its frequent updates can flood the network and its slow convergence can be a problem.
- **Border Gateway Protocol (BGP)** has been the core Internet backbone routing protocol since the mid-1990s, and is the most used interior gateway protocol (IGP). However, some configurations require full mesh connections which flood the network, and there can be route flap and load balancing issues for multihomed networks.
- **Open Shortest Path First (OSPF)** is commonly used in large enterprise networks. It is the protocol of choice, mainly due to its fast convergence. However, it can be complicated to setup properly.
- **Intermediate System to Intermediate System (IS-IS) Protocol** allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) not intended to be used between Autonomous Systems (ASes). IS-IS is a link state protocol well-suited to smaller networks that is in widespread use and has near universal support on routing hardware.
- **Multicast** addressing is used to broadcast from one source to many destinations efficiently. Protocol Independent Multicast (PIM) is the protocol commonly used in enterprises, multimedia content delivery, and stock exchanges.

Routing algorithm

Each protocol uses a slightly different algorithm for choosing the best route between two addresses on the network. The algorithm is the “intelligent” part of a dynamic protocol because the algorithm is responsible for deciding which route is best and should be added to the local routing table. RIP and BGP use distance vector algorithms, where OSPF and IS-IS use link-state or a shortest path first algorithm.

Vector algorithms are essentially based on the number of hops between the originator and the destination in a route, possibly weighting hops based on how reliable, fast, and error-free they are.

The link-state algorithm used by OSPF and IS-IS is called the Dijkstra algorithm. Link-state treats each interface as a link and records information about the state of the interface. The Dijkstra algorithm creates trees to find the shortest paths to the routes it needs based on the total cost of the parts of the routes in the tree.

For more information about the routing algorithm used, see Comparison of dynamic routing protocols on page 107.

Authentication

If an attacker gains access to your network, they can masquerade as a router on your network to either gain information about your network or disrupt network traffic. If you have a high quality firewall configured, it will help your network security and stop many of these types of threats. However, the main method for protecting your routing information is to use authentication in your routing protocol. Using authentication on your FortiGate unit and other routers prevents access by attackers because all routers must authenticate with passwords, such as MD5 hash passwords, to ensure they are legitimate routers.

When you configure authentication on your network, ensure that you configure it the same way on all devices on the network. Failure to do so will create errors and outages as those forgotten devices fail to connect to the rest of the network.

For example, to configure an MD5 key of `123` on an OSPF interface called `ospf_test`, enter the following CLI commands:

```
config router ospf
   config ospf-interface
      edit ospf_test
         set authentication md5
         set md5-key 123
      end
   end
```

Convergence

Convergence is the ability of a networking protocol to re-route around network outages. Static routing cannot do this. Dynamic routing protocols can all converge, but take various amounts of time to do this. Slow convergence can cause problems, such as network loops, which degrade network performance.

You may also hear robustness and redundancy used to describe networking protocols. In many ways, they are the same thing as convergence. Robustness is the ability to keep working even though there are problems, including configuration problems as well as network outages. Redundancy involves having duplicate parts that can continue to function in the event of some malfunction, error, or outage. It is relatively easy to configure dynamic routing protocols to have backup routers and configurations that will continue to function no matter the network problem, short of a total network failure.

IPv6 support

IPv4 addressing is in common use everywhere around the world. IPv6 has much larger addresses and it is used by many large companies and government departments. IPv6 is not as common as IPv4 yet, but more companies are adopting it.

If your network uses IPv6, your dynamic routing protocol must support it. None of the dynamic routing protocols supported IPv6 originally, but they all have additions, expansions, or new versions that now support IPv6. For more information, see RIP on page 120, BGP on page 199, OSPF on page 158, or IS-IS on page 238.

When to adopt dynamic routing

Static routing is more than enough to meet your networking needs when you have a small network. However, as your network grows, the question you need to answer is at what point do you adopt dynamic routing in your networking plan and start using it in your network? The main factors in this decision are typically:

- Budget
- Current network size and topology
- Expected network growth
- Available resources for ongoing maintenance

Budget

When making any business decision, you must always consider your budget. Static routing does not involve special hardware, fancy software, or expensive training courses.

Dynamic routing can include all of these extra expenses. Any new hardware, such as routers and switches, will need to support the routing protocols that you choose. Network management software and routing protocol drivers may also be necessary to help configure and maintain your more complex network. If the network administrators are not well versed in dynamic routing, you must budget either a training course or some hands-on learning time so they can administer the new network with confidence. Together, these factors can impact your budget.

Additionally, people will always account for network starting costs in the budgets but usually leave out the ongoing cost of network maintenance. Any budget must provide for the hours that will be spent on updating the network routing equipment and fixing any problems. Without that money in the budget, you may end up back at static routing before you know it.

Current network size and topology

As stated earlier, static routing works well on small networks. As those networks get larger, routing takes longer, routing tables get very large, and general performance is not what it could be.

Topology is a concern as well. If all your computers are in one building, it is much easier to stay with static routing longer. However, connecting a number of locations will be easier with the move to dynamic routing.

If you have a network of 20 computers, you can still likely use static routing. If those computers are in two or three locations, static routing will still be a good choice for connecting them. Also, if you just connect to your ISP and do not worry about any special routing to do that, you are likely safe with just static routing.

If you have a network of 100 computers in one location, you can use static routing but it will be slower, more complex, and there will not be much room for expansion. If those 100 computers are spread across three or more locations, dynamic routing is the way to go.

If you have 1000 computers, you definitely need to use dynamic routing no matter how many locations you have.

Hopefully this section has given you an idea of what results you will likely experience from different sized networks using different routing protocols. Your choice of which dynamic routing protocol to use is partly determined by the network size and topology.

Expected network growth

You may not be sure if your current network is ready for dynamic routing. However, if you are expecting rapid growth in the near future, it is a good idea to start planning for that growth now so you are ready for the coming expansion.

Static routing is very labor intensive. Each network device's routing table needs to be configured and maintained manually. If there is a large number of new computers being added to the network, they each need to have the static routing table configured and maintained. If devices are being moved around the network frequently, they must also be updated each time.

Instead, consider putting dynamic routing in place before the new computers are installed on the network. The installation issues can be worked out with a smaller and less complex network, and when the new computers or routers are added to the network there will be nowhere near the level of manual configuration required. Depending on the level of growth, the labor savings can be significant. For example, in an emergency you can drop a new router into a network or AS, wait for it to receive the routing updates from its neighbors, and then remove one of the neighbors. While the routes will not be the most effective possible, this method is much less work than static routing in the same situation, with less chance of mistakes.

Also, as your network grows and you add more routers, the new routers can help share the load in most dynamic routing configurations. For example, if you have 4 OSPF routers and 20,000 external routes, those few routers will be overwhelmed. But a network with 15 OSPF routers will be better able to handle that number of routes. However, be aware that adding more routers to your network will increase the amount of updates sent between the routers, which will use up a greater part of your bandwidth and use more bandwidth overall.

Available resources for ongoing maintenance

As explained in the budget section, there must be resources dedicated to ongoing network maintenance, upgrades, and troubleshooting. These resources include administrator hours to configure and maintain the network, training for the administrator (if needed), extra hardware and software as needed, and possible extra staff to help the administrator in emergencies. Without these resources, you will quickly find the network reverting to static routing out of necessity. This is because:

- Routing software updates will require time
- Routing hardware updates will require time
- Office reorganizations or significant personnel movement will require time from a networking point of view
- Networking problems that occur, such as failed hardware, require time to locate and fix the problem

If resources to accomplish these tasks are not budgeted, the tasks will either not happen at the required level to continue operation or not happen at all. This will result in both the network administration staff and the network users being very frustrated.

A lack of a maintenance budget will also result in an increasingly heavy reliance on static routing as the network administrators are forced to use quick fixes for problems that come up. This invariably involves going to static routing, and dropping the more complex and time-consuming dynamic routing.

Choosing a routing protocol

One of that hardest decisions in routing can be choosing which routing protocol to use on your network. It can be easy to decide when static routing will not meet your needs, but how can you tell which dynamic routing protocol is best for your network and situation?

Here is a brief look at the routing protocols, including their strongest and weakest points. The steps to choosing your routing protocol are:

1. Answer questions about your network
2. Evaluate your chosen protocol
3. Implement your dynamic routing protocol

Answer questions about your network

Before you can decide what is best for your situation, you need to examine the details of your situation, such as what you have for budget, equipment, and users.

The following questions will help you form a clear idea of your routing needs:

How many computers or devices are on your network?

The number of computers or devices that you have on your network, and whether the devices are all in one location or distributed, matters. All routing protocols can be run on networks of any size, but it can be inefficient to run some routing protocols on very small networks. Also, routers and network hardware that support dynamic routing can be more expensive than more generic routers for static routing.

What applications typically run over the network?

Finding out what applications your users are running will help you determine their needs and the needs of the network regarding bandwidth, quality of service, and other such issues.

What level of service do the users expect from the network?

Different users have different expectations of the network. It's not critical for someone surfing the Internet to have 100% uptime, but it is required for a stock exchange network or a hospital.

Is there network expansion in your near future?

You may have a small network now, but if it will be growing quickly, you should plan for the expected size so you do not have to change technologies again down the road.

What routing protocols do your networks connect to?

This is most often how routing protocol decisions are made. You need to be able to communicate easily with your service provider and neighbors, so often people simply use what everyone else is using.

Is security a major concern?

Some routing protocols have levels of authentication and other security features built in, and others do not. If security is important to you, be aware of this.

What is your budget?

You need to know what both your initial and maintenance budget is. More robust and feature laden routing protocols generally mean more resources are required to keep them working well. Also, more secure configurations require still more resources. This includes both set up costs and ongoing maintenance costs. If you ignore these costs, you risk having to drop the adoption of the new routing protocol mid-change.

Evaluate your chosen protocol

Once you have examined the features of the routing protocols listed above and chosen the one that best meets your needs, you can set up an evaluation or test installation of that protocol.

The test installation is generally set up in a sandbox configuration so it will not affect critical network traffic. The aim of the test installation is to prove that it will work on a larger scale on your network. You must ensure that the test installation mirrors your larger network well enough for you to discover any problems. If the test installation is too simpe, these problems may not appear.

If your chosen protocol does not meet your goals, choose a different protocol and repeat the evaluation process until a protocol meets your needs or you change your criteria.

Implement your dynamic routing protocol

You have examined your needs, selected the best matching dynamic routing protocol, tested it, and now you are ready to implement it with confidence.

This guide will help you configure your FortiGate unit to support your chosen dynamic routing protocol. Refer to the various sections in this guide, as needed, during your implementation to help ensure a smooth transition. Examples for each protocol are included to show proper configurations for different types of networks.

Dynamic routing terminology

Dynamic routing is a complex subject. There are many routers on different networks and all can be configured differently. It is more complicated by the fact that each routing protocol has different names for similar features, as well as many features that you can configure for each protocol.

To better understand dynamic routing, the following sections provide explanations on common dynamic routing terms

For more details about a term, as it applies to a dynamic routing protocol, see BGP on page 199, RIP on page 120, or OSPF on page 158.

Aggregated routes and addresses

Just as an aggregate interface combines multiple interfaces into one virtual interface, an aggregate route combines multiple routes into one route. This reduces the amount of space those routes require in the routing tables of the routers along that route. The trade-off is a small amount of processing to aggregate and de-aggregate the routes at either end.

The benefit of this method is that you can combine many addresses into one, potentially reducing the routing table size immensely. The weakness of this method is if there are holes in the address range you are aggregating, you need to decide if it is better to break it into multiple ranges, or accept the possibility of failed routes to the missing addresses.

For information about aggregated routes in BGP, see BGP on page 199.

To manually aggregate the range of IP addresses from 192.168.1.100 to 192.168.1.103

1. Convert the addresses to binary:
```
192.168.1.100 = 11000000 10101000 00000001 01100100
192.168.1.101 = 11000000 10101000 00000001 01100101
192.168.1.102 = 11000000 10101000 00000001 01100110
192.168.1.103 = 11000000 10101000 00000001 01100111
```
2. Determine the maximum number of matching bits common to the addresses.
   There are 30-bits in common, with only the last 2-bits being different.
3. Record the common part of the address:
```
11000000 10101000 00000001 0110010X = 192.168.1.100
```
4. For the netmask, assume all the bits in the netmask are 1, except those that are different (which are 0):
```
11111111 11111111 11111111 11111100 = 255.255.255.252
```
5. Combine the common address bits and the netmask:
```
192.168.1.100/255.255.255.252
```
   Alternately, the IP mask may be written as a single number:
```
192.168.1.100/2
```
6. As required, set variables and attributes to declare that the routes have been aggregated, and which router did the aggregating.

Autonomous system

An Autonomous System (AS) is one or more connected networks that use the same routing protocol, and appear to be a single unit to any externally connected networks. For example, an ISP may have a number of customer networks connected to it, but to any networks connected externally to the ISP, it appears as one system or AS. An AS may also be referred to as a routing domain.

It should be noted that while OSPF routing takes place within one AS, the only part of OSPF that deals with the AS is the AS border router (ASBR).

There are multiple types of ASs, which are defined by how they are connected to other ASs. A multihomed AS is connected to at least two other ASs and has the benefit of redundancy: if one of those ASs goes down, your AS can still reach the Internet through its other connection. A stub AS has only one connection, and can be useful in specific configurations where limited access is desirable.

Each AS has a number assigned to it, known as an ASN. In an internal network, you can assign any ASN you like (a private AS number), but for networks connected to the Internet (public AS), you need to have an officially registered ASN from the Internet Assigned Numbers Authority (IANA). ASNs from 1 to 64,511 are designated for public use.

NAs of January 2010, AS numbers are 4 bytes long, instead of the former 2 bytes. RFC 4893 introduced 32-bit ASNs, which FortiGate units support for BGP and OSPF.

Do you need your own AS?

The main factors in deciding if you need your own AS, or if you should be part of someone else's are:

- Exchanging external routing information
- Many prefixes should exist in one AS as long as they use the same routing policy
- When you use a different routing protocol than your border gateway peers. For example, your ISP uses BGP and you use OSPF.
- Connected to many other ASs (multi-homed)

You should not create an AS for each prefix on your network. You also should not be forced into an AS just so someone else can make AS-based policy decisions on your traffic.

There can be only one AS for any prefix on the Internet. This is to prevent routing issues.

What AS number should you use?

In addition to overseeing IP address allocation and Domain Name Systems (DNS), the Internet Assigned Numbers Authority (IANA) assigns public AS numbers. The public AS numbers range from 1 to 64,511. The ASNs 0, 54272 to 64511, and 65535 are reserved by the IANA and should not be used.

ASNs are assigned in blocks by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIR), who then assign ASNs to companies within the geographic area of the RIR. These companies are usually ISPs, and to receive an ASN you must complete the application process of the local RIR and be approved before being assigned an ASN. The following table shows the names and regions of the RIRs:

AFRINIC	Serves the African continent
APNIC	Asia-Pacific, including China, India, and Japan
ARIN	American registry, including Canada and United States
LACNIC	Latin America, including Mexico, Caribbean, Central and South America
RIPE NCC	Europe, the Middle East,the former USSR, and parts of Central Asia

AS numbers from 64512 to 65534 are reserved for private use. Private AS numbers can be used for any internal networks with no outside connections to the Internet, such as test networks, classroom labs, and other internal-only networks that do not access the outside world. You can also configure border routers to filter out any private ASNs before routing traffic to the outside world. If you must use private ASNs with public networks, this is the only way to configure them. However, it is risky because many other private networks could be using the same ASNs and conflicts couldl happen. It would be like your local 192.168.0.0 network being made public and the resulting problems would be widespread.

In 1996, when RFC 1930 was written, only 5,100 ASs had been allocated and a little under 600 ASs were actively routed in the global Internet. Since that time, many more public ASNs have been assigned, leaving only a small number. For this reason 32-bit ASNs (four-octet ASNs) were defined to provide more public ASNs. RFC 4893 defines 32-bit ASNs, and FortiGate units support these larger ASNs.

Area border router

Routers within an AS advertise updates internally and only to each other. However, routers on the edge of the AS must communicate both with routers inside their AS and routers external to their AS, which are often running a different routing protocol. These routers are called Area Border Routers (ABRs) or edge routers. ABRs often run multiple routing protocols in order to redistribute traffic between different ASs that are running different protocols, such as the edge between an ISP's IS-IS routing network and a large company's OSPF network.

OSPF defines ABRs differently from other routers. In OSPF, an ABR is an OSPF router that connects another AS to the backbone AS, and is a member of all the areas it connects to. An OSPF ABR maintains an LSA database for each area that it is connected to. The concept of the edge router is present, but it is the edge of the backbone instead of the edge of the OSPF supported ASs.

Neighbor routers

Routing involves routers communicating with each other. To do this, routers need to know information about each other. These routers are called neighbor routers and are configured in each routing protocol. Each neighbor has custom settings since some routers may have functionality that other routers lack. Neighbor routers are sometimes called peers.

Generally, neighbor routers must be configured and discovered by the rest of the network before they can be integrated into the routing calculations. This is a combination of the network administrator configuring the new router with its neighbor router addresses, and the routing network discovering the new router, such as the hello packets in OSPF. That discovery initiates communication between the new router and the rest of the network.

Route maps

Route maps are a way for the FortiGate unit to evaluate optimum routes for forwarding packets or suppressing the routing of packets to particular destinations. Compared to access lists, route maps support enhanced packet-matching criteria. In addition, route maps can be configured to permit or deny the addition of routes to the FortiGate unit routing table and make changes to routing information dynamically as defined through route-map rules.

Route maps can be used for limiting both received route updates and sent route updates. This can include the redistribution of routes learned from other types of routing. For example, if you do not want to advertise local static routes to external networks, you could use a route map to accomplish this.

The FortiGate unit compares the rules in a route map to the attributes of a route. The rules are examined in ascending order until one or more of the rules in the route map are found to match one or more of the route attributes.

As an administrator, route maps allow you to group a set of addresses together and assign them a meaningful name. During your configuration, you can use these route-maps to speed up configuration. The meaningful names also ensure that fewer mistakes are made during configuration.

The default rule in the route map, which the FortiGate unit applies last, denies all routes. For a route map to take effect, it must be called by a FortiGate unit routing process.

The syntax for route maps are:

```
config router route-map
   edit <route_map_name>
      set comments
      config rule
         edit <route_map_rule_id>
            set action
            set match-*
            set set-*
            ...
         end
```

The `match-*` commands allow you to match various parts of a route. The `set-*` commands allow you to set routing information once a route is matched.

For an example of how route maps can be used to create receiving or sending “groups” in routing, see BGP on page 199.

Access lists

Use this command to add, edit, or delete access lists. Access lists are filters used by FortiGate unit routing processes. For an access list to take effect, it must be called by a FortiGate unit routing process (for example, a process that supports RIP or OSPF). Use `access-list6` for IPv6 routing.

Access lists can be used to filter which updates are passed between routers or which routes are redistributed to different networks and routing protocols. You can create lists of rules that will match all routes for a specific router or group of routers.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or match the prefix and a more specific prefix.

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 cannot be exactly matched with an access-list. A prefix-list must be used for this purpose.

The FortiGate unit attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found, the default action is deny.

The syntax for access lists is:

```
config router access-list, access-list6
   edit <access_list_name>
      set comments
      config rule
      edit <access_list_id>
         set action
         set exact-match
         set prefix
         set prefix6
         set wildcard
```

For an example of how access lists can be used to create receiving or sending “groups” in routing, see BGP on page 199.

Bi-directional forwarding detection

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD send packets to each other at a negotiated rate. If packets from a BFD-protected router fail to arrive, then that router is declared to be down. BFD communicates this information to the routing protocol and the routing information is updated.

BFD neighbors establish if BFD is enabled in OSPF, or BFP routers establish as neighbors.

The CLI commands associated with BFD include:

```
config router bgp
   config neighbor
      set bfd
end
config router ospf
   set bfd
```

```
end
```

Per-VDOM configuration:

```
config system settings
   set bfd
   set bfd-desired-min-tx
   set bfd-required-min-rx
   set bfd-detect-mult
   set bfd-dont-enforce-src-port
end
```

Per-interface (override) configuration:

```
config system interface
   edit <interface_name>
      set bfd enable
      set bfd-desired-min-tx
      set bfd-detect-mult
      set bfd-required-min-rx
   end
```

For more information about BFD in BGP, see BGP on page 199.

Controlling how routing changes affect active sessions

Dynamic routing changes can occur while the FortiGate unit is processing traffic. Routing changes that affect the routes being used for current sessions, may affect how the FortiGate continues to process the session. In FortiOS 5.6.1 and later, you can control how active sessions are affected when dynamic routing changes occur that affects the routes the active sessions are using.

You can configure whether the FortiGate maintains the original routing for the sessions that are using the affected routes, or applies the routing table changes to the active sessions, which may cause destinations to change.

Configure how dynamic routing changes affect active sessions

To configure how dynamic routing changes affect active sessions, use the following CLI commands:

```
config system interface
   edit <port#>
      set preserve-session-route {enable | disable}
   end
```

CLI option	Description
<port#1>	The name of the interface where you want to configure how dynamic routing changes affect active sessions running through it.
enable (default)	All sessions passing through the interface when the routing changes occur, are allowed to finish and are not affected by the routing changes.
disable	When a routing change occurs, the new routing table is applied to the active sessions passing through the interface. The routing changes may cause the destinations of the sessions to change.

IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming more popular and new versions of the dynamic routing protocols have been introduced.

Dynamic routing supports IPv6 on your FortiGate unit. The new versions of these protocols and the corresponding RFCs are:

- **RIP next generation (RIPng)** — RFC 2080 - Routing Information Protocol next generation (RIPng). See RIP and IPv6.
- **BGP4+** — RFC 2545, and RFC 2858 Multiprotocol Extensions for IPv6 Inter-Domain Routing, and Multiprotocol Extensions for BGP-4 (MP-BGP) respectively. See BGP and IPv6.
- **OSPFv3** — RFC 2740 Open Shortest Path First version 3 (OSPFv3) for IPv6 support. See OSPFv3 and IPv6.
- **Integrated IS-IS** — RFC 5308 for IPv6 support. See Integrated IS-IS.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in **System > Admin > Settings**. Alternatively, you can directly configure IPv6 for RIP, BGP, or OSPF protocols using CLI commands.

RIP

This section describes the Routing Information Protocol (RIP).

RIP background and concepts

Background

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. Its widespread use started when an early version of RIP was included with BSD v4.3 Linux as the routed daemon. The Bellman–Ford algorithm, which is the routing algorithm used by RIP, first saw widespread use as the initial routing algorithm of the ARPANET.

RIP has many benefits. It is well suited to smaller networks, has near universal support on routing hardware, is quick to configure, works well if there are no redundant paths, and is in widespread use. However, because RIP updates are sent out node-by-node, it can be slow to find a path around network outages. RIP also lacks good authentication, cannot choose routes based on different quality of service methods, and can create network loops if you are not careful.

The FortiGate implementation of RIP supports RIP version 1 (see RFC 1058), RIP version 2 (see RFC 2453), and the IPv6 version RIPng (see RFC 2080).

RIPv1

In 1988, RIP version 1 (RIPv1) was released. It is defined in RFC 1058. It uses classful addressing and uses broadcasting to send out updates to router neighbors. There is no subnet information included in the routing updates in classful routing. It does not support CIDR addressing and subnets must all be the same size. Also, route summarization is not possible. RIPv1 has no router authentication method, so it is vulnerable to attacks through packet sniffing and spoofing.

RIPv2

In 1993, RIP version 2 (RIPv2) was developed to deal with the limitations of RIPv1. It was not standardized until 1998. This new version supports classless routing and subnets of various sizes. Router authentication was added, which supports MD5. MD5 hashes are an older encryption method, but this is much improved over no security at all. In RIPv2, the hop count limit remained at 15, in order to be backwards compatible with RIPv1. It also uses multicasting to send the entire routing table to router neighbors, which reduces the traffic for devices that are not participating in RIP routing. Routing tags were also added, which allow internal routes or redistributed routes to be identified as such.

RIPng

RIPng, defined in RFC 2080, is an extension of RIPv2 and is designed to support IPv6. However, RIPng varies from RIPv2 in that it is not fully backwards compatible with RIPv1. RIPng does not support RIPv1 update authentication and relies on IPsec instead. It does not allow the attaching of tags to routes, as in RIPv2. RIPng requires specific encoding of the next hop for a set of route entries, unlike RIPv2 that encodes the next-hop into each route entry.

RIP terminology and parts

Before you can understand how RIP functions, you need to understand some of the main concepts and parts of RIP.

RIP and IPv6

RIP Next Generation (RIPng) is a new version of RIP and includes support for IPv6.

The FortiGate unit command `config router ripng` is almost the same as `config router rip`, except that IPv6 addresses are used. Also, if you are going to use prefix or access lists with RIPng, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to troubleshoot RIPng, it is the same as with RIP but specify the different protocol and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table or other related information.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ip6-tunnel` to configure the FortiGate unit to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command is not supported in transparent mode.

For example, if you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01::, where it will need access to an IPv4 network again, use the following commands:

```
config system ipv6-tunnel
   edit test_tunnel
      set destination 2002:A0A:A01::
      set interface port1
      set source 2002:C0A8:3201::
   end
end
```

The CLI commands associated with RIPng include:

```
config router ripng
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

Default information originate option

The default information originate option is the second advanced option for RIP in the web-based manager, right after metric. Enabling default-information-originate will generate and advertise a default route into the FortiGate unit's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. RIP does not create the default route unless you use the always option.

Select **Disable** if you experience any issues or if you wish to advertise your own static routes into RIP updates.

You can enable or disable default-information-originate in **Router > Dynamic > RIP**, under **Advanced Options**, or use the CLI.

The CLI commands associated with default information originate include:

```
config router rip
```

```
    set default-information-originate
end
```

Update, timeout, and garbage timers

RIP uses various timers to regulate its performance including an update timer, a timeout timer, and a garbage timer. The FortiGate unit's default timer settings (30, 180, and 120 seconds) are effective in most configurations. If you change these settings, ensure that the new settings are compatible with local routers and access servers.

The timeout period should be at least three times longer than the update period. If the update timer is smaller than the timeout or garbage timers, you will experience an error.

You can set the three RIP timers in **Router > Dynamic > RIP**, under **Advanced Options**, or use the CLI.

The CLI commands associated with garbage, timeout, and update timers include:

```
config router rip
   set timeout-timer
   set update-timer
   set garbage-timer
end
```

Update timer

The update timer determines the interval between routing updates. This value is usually set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers attempting to update their neighbors simultaneously. The update timer should be at least three times smaller than the timeout timer or you will experience an error.

If you are experiencing significant RIP traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience timeouts that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the FortiGate unit will keep a reachable route in the routing table while no updates for that route are received. If the FortiGate unit receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period or you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods. It may be a considerable amount of time before the FortiGate unit is done waiting for all the timers to expire on unresponsive routes.

Garbage timer

The garbage timer is the amount of time (in seconds) that the FortiGate unit will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove older ones faster. This will result in a smaller routing table, which is useful if you have a very large network, or if your network changes frequently.

Authentication and key chain

RIP version 2 (RIPv2) uses authentication keys to ensure that the routing information exchanged between routers is reliable. RIP version 1 (RIPv1) has no authentication. For authentication to work, both the sending and receiving routers must be set to use authentication and must be configured with the same keys.

The sending and receiving routers need to have their system dates and times synchronized to ensure both ends are using the same keys at the proper times. However, you can overlap the key lifetimes to ensure that a key is always available even if there is some difference in the system times.

A key chain is a list of one or more authentication keys, including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. The FortiGate unit migrates from one key to the next according to the scheduled send and receive lifetimes.

The key-chain command is a CLI router command. You use this command to manage RIPv2 authentication keys. You can add, edit, or delete keys identified by the specified key number.

This example shows how to configure a key chain with two keys that are valid sequentially in time. This example creates a key chain called “rip_key” that has a password of “fortinet”. The accepted and send lifetimes are both set to the same values: a start time of 9:00 am on February 23, 2010 and an end time of 9:00 am on March 17, 2010. A second key is configured with a password of “my_fortigate” that is valid from March 17, 2010 9:01am to April 1 2010 9:00am. This “rip_key” key chain is then used on the port1 interface in RIP.

```
config router key-chain
   edit "rip_key"
      config key
      edit 1
         set accept-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
         set key-string "fortinet"
         set send-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
      next
      edit 2
         set accept-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
         set key-string "my_fortigate"
         set send-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
      next
      end
   end
config router rip
   config interface
      edit port1
         set auth-keychain "rip_key"
      end
   end
```

The CLI commands associated with authentication keys include:

```
config router key-chain

config router rip
   config interface
      edit <interface>
         set auth-keychain
         set auth-mode
         set auth-string
      end
   end
```

Access lists

Access lists are filters used by the FortiGate unit's RIP and OSPF routing. An access list provides a list of IP addresses and the action to take for them. Essentially, an access list makes it easy to group addresses that will be treated the same way into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include and specify the action to take for it. For example, if you want all traffic from one department to be routed a particular way, even in different buildings, you can add all of the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found, the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. Also, it is easier to troubleshoot because if all addresses on one list have problems, many possible causes can be eliminated right away.

If you are using the RIPng or OSPF+ IPv6 protocols, you will need to use access-list6, which is the IPv6 version of the access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of `10.10.10.10` and `11.11.11.11`, enter the command:

```
config router access-list
   edit test_list
      config rule
         edit 1
            set prefix 10.10.10.10 255.255.255.255
            set action allow
            set exact-match enable
         next
         edit 2
            set prefix 11.11.11.11 255.255.255.255
            set action allow
            set exact-match enable
         end
      end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the access-list6 command, as follows:

```
config router access-list6
   edit test_list_ip6
      config rule
         edit 1
            set prefix6 2002:A0A:A0A:0:0:0:0:0:/48
            set action deny
         next
         edit 2
            set prefix6 2002:B0B:B0B:0:0:0:0:0/48
            set action deny
         end
      end
```

To use an access list, you must call it from a routing protocol, such as RIP. The following example uses the access list from the previous example, called test_list, to match routes coming in on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially increase. Enter the following command:

```
config router rip
   config offset-list
      edit 5
         set access-list test_list
         set direction in
         set interface port1
         set offset 3
         set status enable
      end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 cannot be exactly matched with an access list. A prefix list must be used for this purpose

How RIP works

As one of the original modern dynamic routing protocols, RIP is straightforward. Its routing algorithm is not complex and there are some options that allow fine tuning. It is relatively simple to configure RIP on FortiGate units.

From RFC 1058:

> *Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.*

This section includes:

- RIP versus static routing
- RIP hop count
- The Bellman–Ford routing algorithm
- Passive versus active RIP interfaces
- RIP packet structure

RIP versus static routing

RIP was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, RIP is a big step forward from simple static routing.

While RIP may be slow in response to network outages, static routing has zero response. The same is true for convergence; static routing has zero convergence. Both RIP and static routing have the limited hop count, so it is not a strength or a weakness. Count to infinity can be a problem, but can typically be fixed as it happens, or is the result of a network outage that would cause even worse problems on a static routing network.

This compares to static routing where each time a packet needs to be routed, the FortiGate unit can send it only to the next hop towards the destination. That next hop then forwards it, and so on until it arrives at its destination. RIP keeps more routing information on each router so your FortiGate unit can send the packet further towards its destination before it has to be routed again toward its destination. RIP uses a smaller amount of table lookups,

and therefore fewer network resources, than static routing. Also, since RIP is updated on neighboring routes, it is aware of new routes or dead routes that static routing would not be aware of.

Overall, RIP is a large step forward when compared to static routing.

RIP hop count

RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to the FortiGate unit, while a hop count of 16 represents a network that cannot be reached. Each network that a packet travels through to reach its destination usually counts as one hop. When the FortiGate unit compares two routes to the same destination, it adds the route having the lowest hop count to the routing table. As you can see in RIP packet structure on page 130, the hop count is part of a RIP v2 packet.

Similarly, when RIP is enabled on an interface, the FortiGate unit sends RIP responses to neighboring routers on a regular basis. The updates provide information about the routes in the FortiGate unit's routing table, subject to the rules that you specify for advertising those routes. You can specify how often the FortiGate unit sends updates, the period of time a route can be kept in the routing table without being updated, and for routes that are not updated regularly, you can specify the period of time that the unit advertises a route as unreachable before it is removed from the routing table.

If hops are weighted higher than one, it is very easy to reach the upper limit. This higher weighting will effectively limit the size of your network, depending on the numbers used. Merely changing from the default of 1.0 to 1.5 will lower the effective hop count from 15 to 10. This is acceptable for smaller networks, but can be a problem as your network expands over time.

In RIP, you can use the offset command to artificially increase the hop count of a route. Doing this will make this route less preferred and, in turn, it will get less traffic. Offsetting routes is useful when you have network connections that have different bandwidths, levels of reliability, or costs. In each of these situations you still want the redundancy of multiple route access, but you do not want the bulk of your traffic using these less preferred routes. For an example of RIP offset, see Access lists on page 124.

The Bellman–Ford routing algorithm

The routing algorithm used by RIP was first used in 1967 as the initial routing algorithm of the ARPANET. The Bellman–Ford algorithm is distributed because it involves a number of nodes (routers) within an Autonomous system, and consists of the following steps:

1. Each node calculates the distances between itself and all other nodes within the AS and stores this information as a table.
2. Each node sends its table to all neighboring nodes.
3. When a node receives distance tables from its neighbors, it calculates the shortest routes to all other nodes and updates its own table to reflect any changes.

To examine how this algorithm functions let us look at a network with 4 routers: routers 1 through 4. The distance from Router1 to Router2 is 2 hops, Router1 to Router3 is 3 hops, and Router2 to Router3 is 4 hops. Router4 is only connected to Router2 and Router3, each distance being 2 hops.

1. Router1 finds all of the distances to the other three routers: Router 2 is 2, Router 3 is 3. Router1 does not have a route to Router4.
2. Router2, Router3, and Router4 perform the same calculations from their point of views.
3. Once Router1 gets an update from Router 2 or Router3, it will get their route to Router4. At that point, it now has a route to Router4 and installs that in its local table.

4. If Router1 gets an update from Router3 first, it has a hop count of 5 to reach Router4, but when Router2 sends its update, Router1 will go with Router2's shorter 4 hops to reach Router4. Future updates do not change this unless they are shorter than 4 hops or the routing table route goes down.

RIP algorithm example in four steps

Step 1

Router1 finds the distance to other routers in the network.

It currently has no route to Router4.

Router1 routing table:

- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops

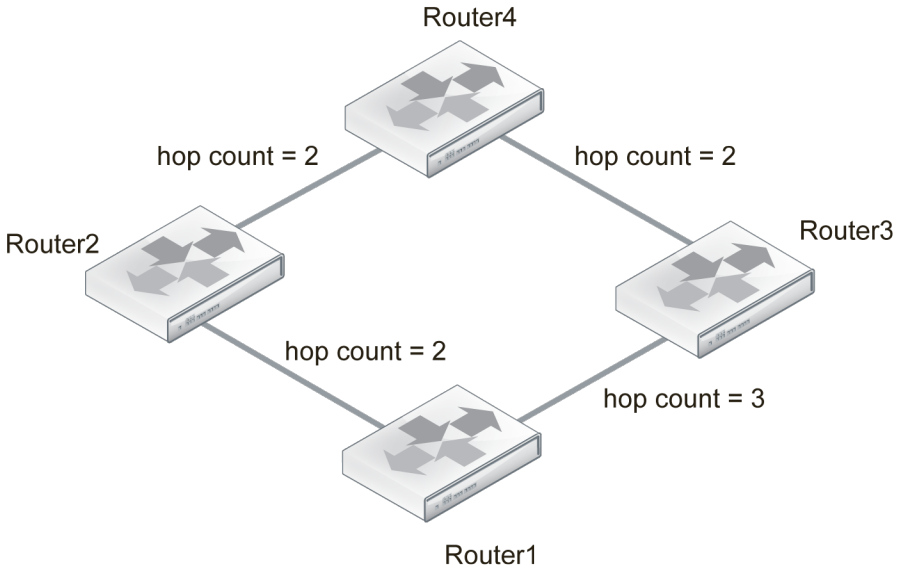

Step 2

All routers do the same as Router1 and send out updates containing their routing table.

Note that Router1 and Router4 do not update each other, but rely on Router2 and Router3 to pass along accurate updates.

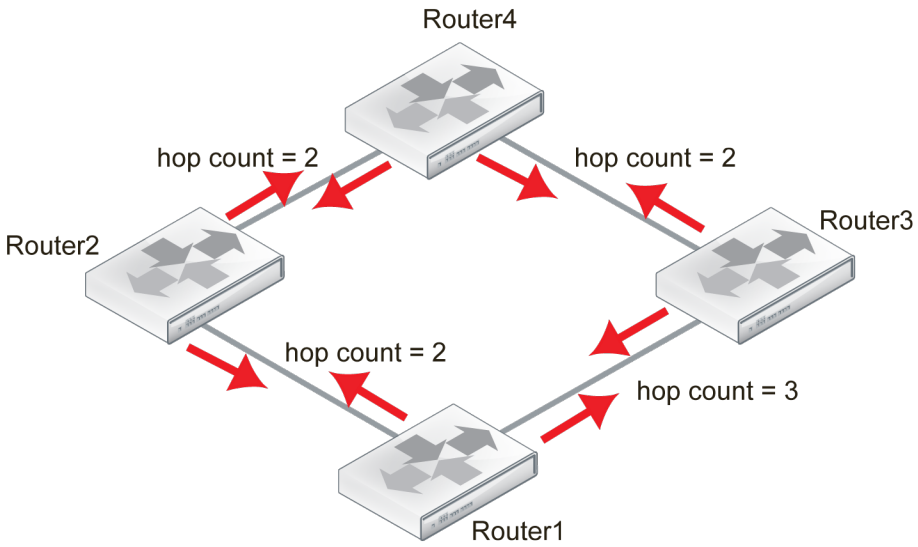

Step 3

Each router looks at the updates it has received and adds any new or shorter routes to its table.

Router1's updated table:

- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops
- Distance to Router4 = 4 or 5 hops

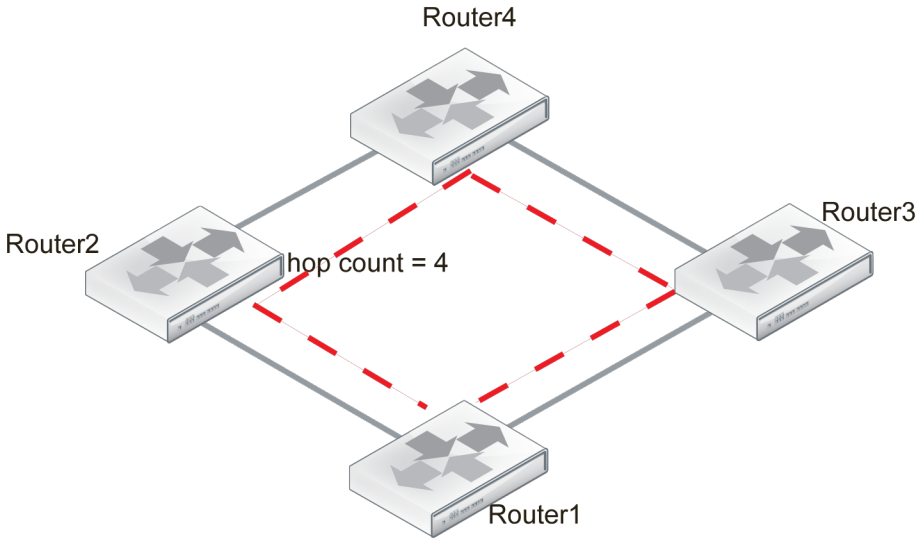

Step 4

Router1 installs the shortest route to Router4 and the other routes to it are removed from the routing table.

Router1's complete table:

- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops

- Distance to Router4 = 4 hops

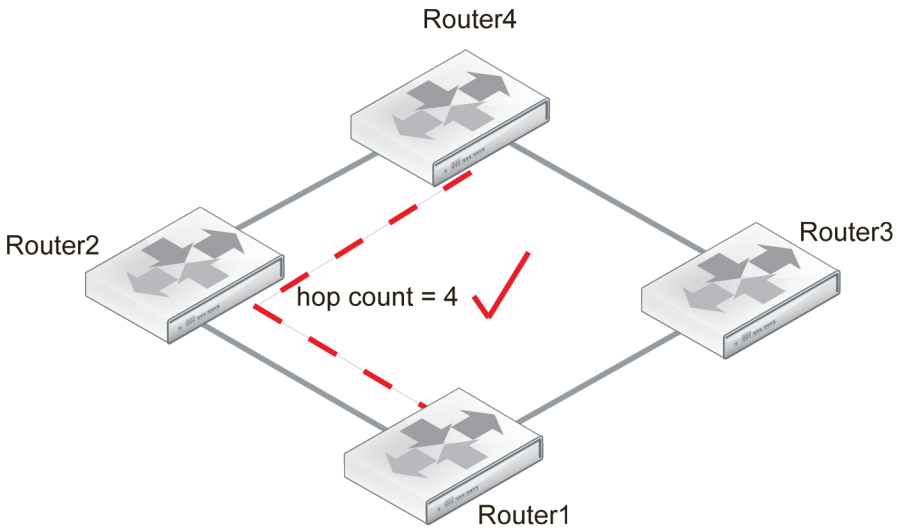

The good part about the Bellman-Ford algorithm in RIP is that the router only uses the information it needs from the update. If there are no newer, better routes than the ones the router already has in its routing table, there is no need to change its routing table. And no change means no additional update, and therefore less traffic. But even when there is update traffic, the RIP packets are very small so it takes many updates to affect overall network bandwidth. For more information about RIP packets, see RIP packet structure on page 130.

The main disadvantage of the Bellman–Ford algorithm in RIP is that it does not take weightings into consideration. While it is possible to assign different weights to routes in RIP, doing so severely limits the effective network size by reducing the hop count limit. Also, other dynamic routing protocols can take route qualities, such as reliability or delay, into consideration to provide not only the physically shortest routes but also the fastest or more reliable routes.

Another disadvantage of the Bellman-Ford algorithm is due to the slow updates passed from one RIP router to the next. This results in a slow response to changes in the network topology, which in turn results in more attempts to use routes that are down and that wastes time and network resources.

Passive versus active RIP interfaces

Normally, the FortiGate unit's routing table is kept up to date by periodically asking the neighbors for routes, and sending your routing updates out. This has the downside of generating a lot of extra traffic for large networks. The solution to this problem is passive interfaces.

A standard interface that supports RIP is active, by default. It sends and receives updates by actively communicating with its neighbors. A passive RIP interface does not send out updates. It only listens to the updates of other routers. This is useful in reducing network traffic, and if there are redundant routers in the network that will send out essentially the same updates all the time.

The following example shows how to create a passive RIPv2 interface on port1 using MD5 authentication and a key chain called `passiveRIPv2`, which has already been configured. Note that in the CLI, you enable passive by disabling `send-version2-broadcast`.

To create a passive RIP interface - web-based manager

1. Go to **Router > Dynamic > RIP**.
2. Next to **Interfaces**, select **Create**.

3. Select port1 as the **Interface**.
4. Select 2 as both the **Send Version** and **Receive Version**.
5. Select MD5 for **Authentication**.
6. Select the `passiveRIPv2`**Key-chain**.
7. Select **Passive Interface**.
8. Select **OK** to accept this configuration and return to the main RIP display page.

To create a passive RIP v2 interface on port1 using MD5 authentication - CLI

```
config router rip
   config interface
      edit port1
         set send-version2-broadcast disable
         set auth-keychain "passiveRIPv2"
         set auth-mode md5
         set receive-version 2
         set send-version 2
      end
   end
```

RIP packet structure

It is hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how it is exchanged will help you to better understand the RIP protocol and better configure your network for it.

This section provides information about the contents of RIPv1 and RIPv2 packets.

RIP version 1

RIP version 1 (RIPv1), or RIP IP, packets are 24 bytes in length with some empty areas left for future expansion.

RIP IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

A RIPv1 packet contains the following fields:

- **Command**: Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.
- **Version**: Specifies the RIP version used. This field can signal different, potentially incompatible versions.
- **Zero field**: This field defaults to zero and is not used by RFC 1058 RIP.
- **Address-family identifier (AFI)**: Specifies the address family used. RIP is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.
- **IP Address**: Specifies the IP address for the entry.
- **Metric**: This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable, the metric is 16.

RIP version 2

RIP version 2 (RIPv2) has more features than RIPv1, which is reflected in its packets that carry more information. All but one of the empty zero fields in RIPv1 packets are used in RIPv2.

RIPv2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag
4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric	

A RIPv2 packet contains the fields described above for RIPv1, as well as the following:

- **Unused**: Has a value set to zero and is intended for future use
- **Route tag**: Provides a method for distinguishing between internal routes learned by RIP and external routes learned from other protocols.
- **Subnet mask**: Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.
- **Next hop**: Indicates the IP address of the next hop to which packets for the entry should be forwarded.

Troubleshooting RIP

This section is about troubleshooting RIP. For general troubleshooting information, see the FortiOS Handbook Troubleshooting chapter.

Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself, creating a loop. When there are loops, the network has problems getting information to its destination. Loops also prevent the network from returning to the source to report the inaccessible destination.

A routing loop occurs when a normally functioning network has an outage and one or more routers are offline. When packets encounter this, they attempt an alternate route maneuver around the outage. During this phase, it is possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is also blocked by the outage, a hop back, and possibly the original hop forward, may be selected. If this continues, it can consume not only network bandwidth but also many resources on the affected routers. The worst part is this situation will continue until the network administrator changes the router settings or the downed routers come back online.

Effect of routing loops on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes, that router sends an update out to all of the RIP routers connected to it. In a network loop, it is possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How to spot a routing loop

Anytime network traffic slows down, you will ask yourself if it is a network loop. Slowdowns are often normal, are not a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown that does not return to normal quickly, you need to do serious troubleshooting quickly.

If you are not running SNMP, dead gateway detection, or you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it. Ping, traceroute, and other basic troubleshooting tools are largely the same between static and dynamic and are covered in Advanced static routing on page 72.

Check your logs

If your routers log events to a central location, it can be easy to check your network logs for any outages.

On your FortiGate unit, go to **Log & Report**. You should look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and its location as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

1. Go to **System > Config > SNMP**.
2. Enable **SMTP Agent** and select **Apply**.
   Optionally, enter the **Description**, **Location**, and **Contact** information for this device for easier location of the problem report.
3. Under **SNMP v1/v2** or **SNMP v3** as appropriate, select **Create New**.
   SNMP v3

User Name	Enter the SNMP user ID.
Security Level	Select authentication or privacy as desired. Select the authentication or privacy algorithms to use and enter the required passwords.
Notification Host	Enter the IP addresses of up to 16 hosts to notify.

Enable Query	Select. The **Port** should be 161. Ensure that your security policies allow ports 161 and 162 (SNMP queries and traps) to pass.
SNMP v1/v2	
Hosts	Enter the IP addresses of up to 8 hosts to notify.
Queries	Enable **v1** and/or **v2** as needed. The **Port** should be 161. Ensure that your security policies allow port 161 to pass.
Traps	Enable v1 and/or v2 as needed. The Port should be 162. Ensure that your security policies allow port 162 to pass.

4. Select the events for which you want notification. For routing loops this should include **CPU usage is high**, **Memory is low**, and possibly **Log disk space is low**. If there are problems the log will fill up quickly and the FortiGate unit's resources will be overused.
5. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Use Link Health Monitor and email alerts

Another tool available to you on FortiGate units is the Link Health Monitor, which is useful for dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with Link Health Monitor and email alerts

Use the following command to configure dead gateway detection:

```
config system link-monitor
   edit "test"
      set srcintf "internal4"
      set server "8.8.8.8"
      set interval 5
      set failtime 1
   end
```

Set the `Interval` (how often to send a ping) and `failtime` (how many lost pings are considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.

To configure notification of failed gateways

1. Go to **Log & Report > Report > Local** and enable **Email Generated Reports**.
2. Enter your email details.
3. Select **Apply**.

You might also want to log CPU and Memory usage as a network outage will cause your CPU activity to spike.

If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email about the outage.

Look at the packet flow

If you want to see what is happening on your network, look at the packets traveling on the network. This is the same idea as police pulling over a car and asking the driver where they have been and what the conditions were like.

The method used in the troubleshooting sections Debugging IPv6 on RIPng on page 135 and on debugging the packet flow also apply here. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable.

Ideally, if you debug the flow of the packets and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network and determined that it is a routing loop, there are a number of steps you can take to correct it:

1. Get any offline routers back online. This may be a simple reboot, or you may have to replace hardware. Often, this first step will restore your network to its normal operation, once the routing tables are updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Holddowns and triggers for updates

One of the potential problems with RIP is the frequent routing table updates that are sent every time there is a change to the routing table. If your network has many RIP routers, these updates can start to slow your network down. Also, if you have a particular route that has bad hardware, it might be going up and down frequently, which will generate an overload of routing table updates.

One of the most common solutions to this problem is to use holddown timers and triggers for updates. These slow down the updates that are sent out and help prevent a potential flood.

Holddown timers

The holddown timer activates when a route is marked down. Until the timer expires, the router does not accept any new information about that route. This is very useful if you have a flapping route because it will prevent your router from sending out updates and being part of the problem in flooding the network. The potential downside is if the route comes back up before the timer expires, that route will be unavailable for that period of time. This is only a problem if this is a major route used by the majority of your traffic. Otherwise, this is a minor problem as traffic can be re-routed around the outage.

Triggers

Triggered RIP is an alternate update structure that is based around limiting updates to only specific circumstances. The most basic difference is that the routing table will only be updated when a specific request is sent to update, instead of every time the routing table changes. Updates are also triggered when a unit is 'powered on', which can include the addition of new interfaces or devices to the routing structure, or devices returning to being available after being unreachable.

Split horizon and poison reverse updates

Split horizon is best explained with an example. If there are three routers linked serially, called routerA, routerB, and routerC. RouterA is only linked to routerB, RouterC is only linked to routerB, and routerB is linked to both routerA and routerC. To get to routerC, routerA must go through routerB. If the link to routerC goes down, it is possible that routerB will try to use routerA's route to get to routerC. This route is A-B-C, so it will loop endlessly between routerA and routerB.

This situation is called a split horizon because from routerB's point of view the horizon stretches out in each direction but in reality it is only on one side. Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This poisoned route is marked as unreachable for routers that cannot use it. In RIP, this means that the route is marked with a distance of 16.

Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit, including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
diagnose ipv6 router rip all enable
```

These three commands will:

- turn on debugging, in general
- set the debug level to information, which is a verbose reporting level
- turn on all RIP router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable, since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a reply received in response.

For more information, see Troubleshooting RIP on page 131.

Simple RIP example

This is an example of a typical medium-sized network configuration using RIP routing.

Your company has 3 small local networks, one for each department. These networks are connected by RIP, and then connected to the Internet. Each subnet has more than one route for redundancy. There are two central routers that are both connected to the Internet and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running RIP, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running, those will need to be redistributed through the RIP network.

To keep the example simple, there will be no authentication of router traffic.

With RIP properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. RIP is good for a smaller network due to its lack of complex configurations.

Network layout and assumptions

Basic network layout

Your company has 3 departments each with their own network: Sales, R&D, and Accounting. Each network has routers that are not running RIP and FortiGate units running RIP.

The R&D network has two RIP routers, and each is connected to both other departments as well as being connected to the Internet through the ISP router. The links to the Internet are indicated in black.

The three internal networks do not run RIP. They use static routing because they are small networks. This means the FortiGate units have to redistribute any static routes they learn so that the internal networks can communicate with each other.

Where possible in this example, the default values will be used (or the most general settings). This is intended to provide an easier configuration that will require less troubleshooting.

In this example, the routers, networks, interfaces used, and IP addresses are as follows. Note that the interfaces that connect Router2 and Router3 also connect to the R&D network.

RIP example network topology

Network	Router	Interface & alias	IP address
Sales	Router1	port1 (internal)	10.11.101.101
		port2 (router2)	10.11.201.101
		port3 (router3)	10.11.202.101

Network	Router	Interface & alias	IP address
R&D	Router2	port1 (internal)	10.12.101.102
		port2 (router1)	10.11.201.102
		port3 (router4)	10.14.201.102
		port4 (ISP)	172.20.120.102
	Router3	port1 (internal)	10.12.101.103
		port2 (router1)	10.11.201.103
		port3 (router4)	10.14.202.103
		port4 (ISP)	172.20.120.103
Accounting	Router4	port1 (internal)	10.14.101.104
		port2 (router2)	10.14.201.104
		port3 (router3)	10.14.202.104

Network topology for the simple RIP example

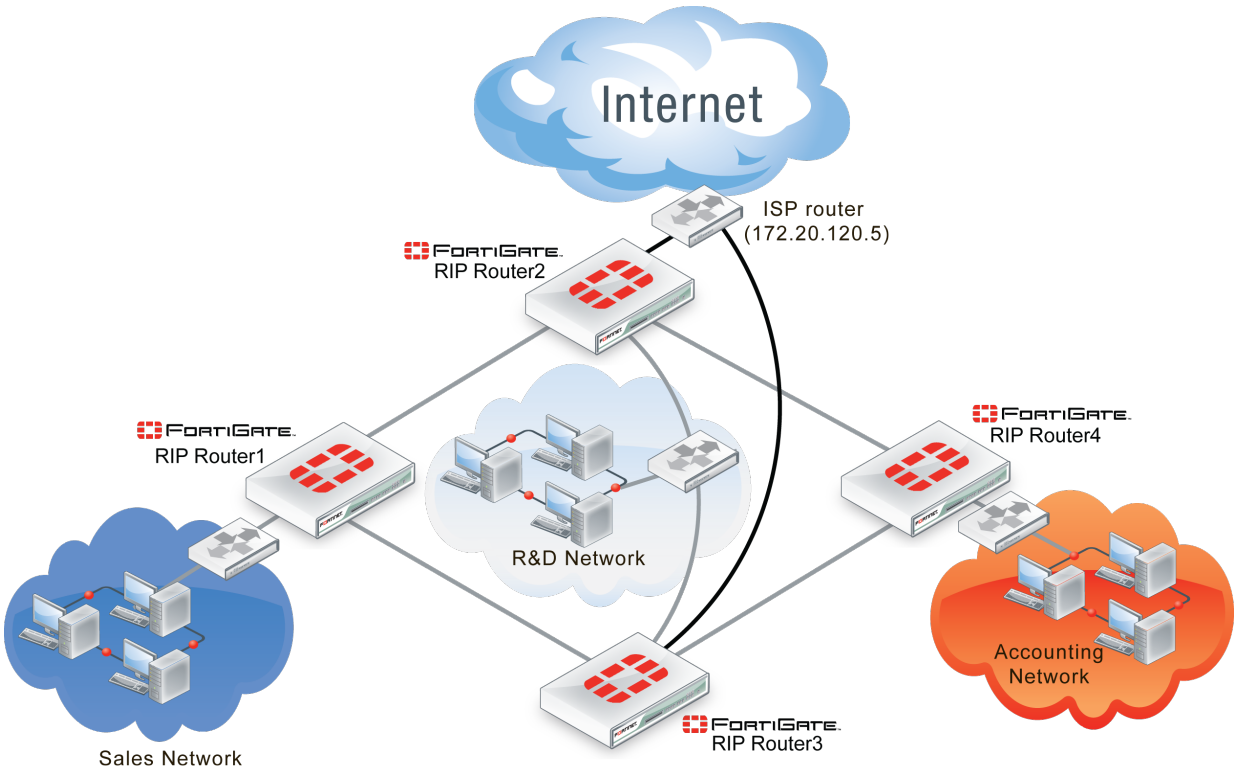

Assumptions

The following assumptions have been made concerning this example:

- All FortiGate units have 5.0 firmware and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labeled port1 through port4, as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- Only FortiGate units are running RIP on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.
- Router2 and Router3 each have their own connection to the Internet, indicated in black in the diagram above.

General configuration steps

This example is very straightforward. The only steps involved are:

- Configuring FortiGate system information
- Configuring FortiGate unit RIP router information
- Configuring other networking devices
- Testing network configuration

Configuring FortiGate system information

You must configure the hostname and interfaces for each FortiGate.

For IP numbering, Router2 and Router3 use the numbering for the other routers, where needed.

Router2 and Router3 have dead gateway detection enabled on the ISP interfaces using Ping. Remember to contact the ISP and confirm their server has ping enabled.

Configure the hostname, interfaces, and default route

To configure Router1 system information - web-based manager

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router1`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port2 (router2)
Administrative Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port3 (router3)
Administrative Distance	40

6. Go to **Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Internal sales network
Interface State	Enabled

9. Edit port2 (router2) interface.
10. Set the following information, and select **OK**.

Alias	router2
IP/Network Mask	10.11.201.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network & Internet through Router2
Interface State	Enabled

11. Edit port3 (router3) interface.
12. Set the following information, and select **OK**.

Alias	router3
IP/Network Mask	10.11.202.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network and Internet through Router3
Interface State	Enabled

To configure Router1 system information - CLI

```
config system global
   set hostname Router1
end

config router static
   edit 1
      set device "port2"
      set distance 45
      set gateway 10.11.201.102
   next
   edit 2
      set device "port3"
      set distance 45
      set gateway 10.11.202.103
   end
end


config system interface
   edit port1
      set alias internal
      set ip 10.11.101.101/255.255.255.0
      set allowaccess https ssh ping
      set description "Internal sales network"
   next
   edit port2
      set alias ISP
      set allowaccess https ssh ping
      set ip 10.11.201.101/255.255.255.0
      set description "Link to R&D network & Internet through Router2"
   next
   edit port3
      set alias router3
      set ip 10.11.202.101/255.255.255.0
      set allowaccess https ssh ping
      set description "Link to R&D network & Internet through Router2"
   end
end
```

To configure Router2 system information - web-based manager

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router2`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port4 (ISP)
Administrative Distance	5

5. Go to **Network > Interfaces**.
6. Edit port1 (internal) interface.
7. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.12.101.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	R&D internal network and Router3
Interface State	Enabled

8. Edit port2 (router1) interface.
9. Set the following information and select **OK**.

Alias	router1
IP/Network Mask	10.12.201.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router1 and the Sales network
Interface State	Enabled

10. Edit port3 (router4) interface.
11. Set the following information and select **OK**.

Alias	router4
IP/Network Mask	10.12.301.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router4 and the accounting network
Interface State	Enabled

12. Edit port4 (ISP) interface.
13. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.102/255.255.255.0
Administrative Access	HTTPS SSH PING

Detect and Identify Devices	enable
Comments	Internet through ISP
Interface State	Enabled

To configure Router2 system information - CLI

```
config system global
   set hostname Router2
end
config router static
   edit 1
      set device "port4"
      set distance 5
      set gateway 172.20.130.5
   end
end
config system interface
   edit port1
      set alias internal
      set ip 10.11.101.102/255.255.255.0
      set allowaccess https ssh ping
      set description "Internal RnD network and Router3"
   next
   edit port2
      set alias router1
      set allowaccess https ssh ping
      set ip 10.11.201.102/255.255.255.0
      set description "Link to Router1"
   next
   edit port3
      set alias router3
      set ip 10.14.202.102/255.255.255.0
      set allowaccess https ssh ping
      set description "Link to Router4"
   next
   edit port4
      set alias ISP
      set ip 172.20.120.102/255.255.255.0
      set allowaccess https ssh ping
      set description "ISP and Internet"
   end
end
```

To configure Router3 system information - web-based manager

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router3`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port4 (ISP)
Administrative Distance	5

5. Go to **Network > Interfaces**.
6. Edit port1 (internal) interface.
7. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.12.101.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	R&D internal network and Router2
Interface State	Enabled

8. Edit port2 (router1) interface.
9. Set the following information and select **OK**.

Alias	router1
IP/Network Mask	10.13.201.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router1 and Sales network
Interface State	Enabled

10. Edit port3 (router4) interface.
11. Set the following information and select **OK**.

Alias	router4
IP/Network Mask	10.13.301.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router4 and accounting network
Interface State	Enabled

12. Edit port4 (ISP) interface.
13. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect and Identify Devices	enable
Comments	Internet and ISP
Interface State	Enabled

To configure Router3 system information - CLI

```
config system global
   set hostname Router3
end
config router static
   edit 1
      set device "port4"
      set distance 5
      set gateway 172.20.130.5
   end
end
config system interface
   edit port1
      set alias internal
      set ip 10.12.101.103/255.255.255.0
      set allowaccess https ssh ping
      set description “Internal RnD network and Router2”
   next
   edit port2
      set alias ISP
      set allowaccess https ssh ping
      set ip 10.11.201.103/255.255.255.0
      set description “Link to Router1”
   next
   edit port3
      set alias router3
      set ip 10.14.202.103/255.255.255.0
      set allowaccess https ssh ping
      set description “Link to Router4”
   next
   edit port4
      set alias ISP
      set ip 172.20.120.103/255.255.255.0
      set allowaccess https ssh ping
      set description “ISP and Internet”
   end
end
```

To configure Router4 system information - web-based manager

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router4`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port2 (router2)
Administrative Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port3 (router3)
Administrative Distance	40

6. Go to **Network > Interfaces**.
7. Edit port 1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.14.101.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Internal accounting network
Interface State	Enabled

9. Edit port 2 (router2) interface.
10. Set the following information and select **OK**.

Alias	router2
IP/Network Mask	10.14.201.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network & Internet through Router2

Interface State	Enabled

11. Edit port 3 (router3) interface.
12. Set the following information and select **OK**.

Alias	router3
IP/Network Mask	10.14.301.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network and Internet through Router3
Interface State	Enabled

To configure Router4 system information - CLI

```
config system global
   set hostname Router4
end
config router static
   edit 1
      set device "port2"
      set distance 45
      set gateway 10.14.201.102
   next
   edit 2
      set device "port3"
      set distance 45
      set gateway 10.14.202.103
   end
end
config system interface
   edit port1
      set alias internal
      set ip 10.14.101.104/255.255.255.0
      set allowaccess https ssh ping
      set description "Internal sales network"
   next
   edit port2
      set alias router2
      set allowaccess https ssh ping
      set ip 10.14.201.104/255.255.255.0
      set description "Link to R&D network & Internet through Router2"
   next
   edit port3
      set alias router3
      set ip 10.14.202.104/255.255.255.0
      set allowaccess https ssh ping
      set description "Link to R&D network & Internet through Router2"
   end
end
```

Configuring FortiGate unit RIP router information

With the interfaces configured, RIP can now be configured on the FortiGate units.

For each FortiGate unit, the following steps will be taken:

- Configure RIP version used
- Redistribute static networks
- Add networks serviced by RIP
- Add interfaces that support RIP on the FortiGate unit

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures. Repeat the procedures once for each FortiGate unit.

Configure RIP settings on Router1 and Router4 - web-based manager

1. Go to **Network > RIP**.
2. Select **2** for **Version**.
3. In **Advanced Options**, under **Redistribute** enable **Static**. Leave the other advanced options at their default values.
4. Under **Networks**, add the following networks:
   - 10.11.0.0/255.255.0.0
   - 10.12.0.0/255.255.0.0
   - 10.14.0.0/255.255.0.0
   - 172.20.120.0/255.255.255.0
6. Under **Interfaces**, select **Create New** and set the following information:

Interface	port1 (internal)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

7. Under **Interfaces** select **Create New**and set the following information:

Interface	port2 (router2)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

8. Under **Interfaces**, select **Create New**and set the following information:

Interface	port3 (router3)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

Configure RIP settings on Router1 and Router4 - CLI

```
config router rip
   set version 2
   config interface
      edit "port1"
         set receive-version 1 2
         set send-version 1 2
      next
      edit "port2"
         set receive-version 1 2
         set send-version 1 2
      next
      edit "port3"
         set receive-version 1 2
         set send-version 1 2
      end
   config network
      edit 1
         set prefix 10.11.0.0 255.255.0.0
      next
      edit 2
         set prefix 10.12.0.0 255.255.0.0
      next
      edit 3
         set prefix 10.14.0.0 255.255.0.0
      next
      edit 4
         set prefix 172.20.120.0 255.255.255.0
      end
   config redistribute "static"
      set status enable
   end
end
```

Configure RIP settings on Router2 and Router3 - web-based manager

1. Go to **Network > RIP**.
2. Select **2** for **RIP**.
3. In **Advanced Options**, under **Redistribute** enable **Static**. Leave the other advanced options at their default values.
4. Under **Networks**, add the following networks:

- 10.11.0.0/255.255.0.0
- 10.12.0.0/255.255.0.0
- 10.14.0.0/255.255.0.0
- 172.20.120.0/255.255.255.0

6. Under **Interfaces**, select **Create New** and set the following information:

Interface	port1 (internal)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

7. Under **Interfaces**, select **Create New** and set the following information:

Interface	port2 (router1)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

8. Under **Interfaces**, select **Create New** and set the following information:

Interface	port3 (router4)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

9. Under **Interfaces**, select **Create New** and set the following information:

Interface	port4 (ISP)
Passive	disabled
Authentication	None

Send Version	Both
Receive Version	Both

Configure RIP settings on Router2 and Router3 - web-based manager

```
config router rip
   set version 2
   config interface
      edit "port1"
         set receive-version 1 2
         set send-version 1 2
      next
      edit "port2"
         set receive-version 1 2
         set send-version 1 2
      next
      edit "port3"
         set receive-version 1 2
         set send-version 1 2
      end
      edit "port4"
         set receive-version 1 2
         set send-version 1 2
      end
   config network
      edit 1
         set prefix 10.11.0.0 255.255.0.0
      next
      edit 2
         set prefix 10.12.0.0 255.255.0.0
      next
      edit 3
         set prefix 10.14.0.0 255.255.0.0
      next
      edit 4
         set prefix 172.20.120.0 255.255.255.0
      end
   config redistribute "static"
      set status enable
   end
end
```

Configuring other networking devices

In this example, there are two groups of other devices on the the network: internal devices and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers, and other network devices. Once the FortiGate units are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise, there should be no configuration required.

The second group of devices is the ISP. This consists of the RIP router the FortiGate Router2 and Router3 connect to. You need to contact your ISP and ensure they have your information for your network, such as the IP

addresses of the connecting RIP routers, what version of RIP your network supports, and what authentication (if any) is used.

Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the Internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult Simple RIP example on page 136.

IPsec auto discovery support

The following routing settings are available in the CLI to support IPsec auto discovery. They are designed for:

- Supporting the RIPng (RIP next generation) network command
- Limiting the maximum metric allowed to output for RIPng
- Fix NSM missing kernel address update information

The actual new settings are:

```
config router rip
   set max-out-metric <integer value 1 - 15>
   end

config router ripng
   set max-out-metric <integer value 1 - 15>
   end

config router ripng
   config network
      edit <ID # of network>
         set prefix <IPv6 prefix>
         end
      end
```

RIPng: RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network, allowing it to reach the Internet at all times.

Network layout and assumptions

Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network, allowing it to reach the Internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.

Where possible in this example, the default values will be used (or the most general settings). This is intended to provide an easier configuration that will require less troubleshooting.

In this example, the routers, networks, interfaces used, and IP addresses are as follows.

RIP example network topology

Network	Router	Interface & alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

Network topology for the IPV6 RIPng example

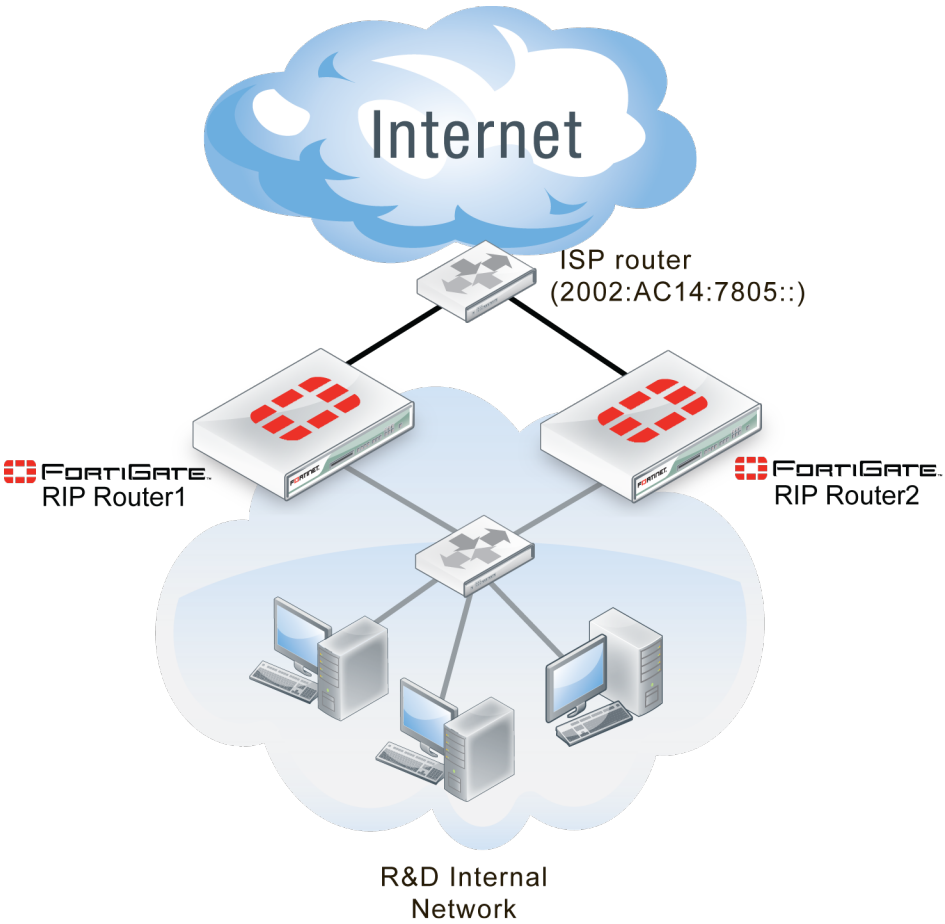

Assumptions

The following assumptions have been made concerning this example.

- All FortiGate units have 5.0 firmware and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labeled port1 and port2, as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices support IPv6 and are running RIPng.

Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.

To configure system information on Router1 - web-based manager

1. Go to **System > Dashboard > Status**.
2. For **Host name**, select **Change**.
3. Enter “Router1”.

4. Go to **System > Config > Features**.
5. In **Basic Features**, enable **IPv6**, and select **Apply**.
6. Go to **System > Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	2002:A0B:6565::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	2002:AC14:7865::/0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet
Administrative Status	Up

To configure system information on Router1 - CLI

```
config system global
   set hostname Router1
   set gui-ipv6 enable
end
config system interface
   edit port1
      set alias internal
      set allowaccess https ping ssh
      set description “Internal RnD network”
      config ipv6
         set ip6-address 2002:a0b:6565::/0
      end
   next
   edit port2
      set alias ISP
      set allowaccess https ping ssh
      set description “ISP and Internet”
      config ipv6
         set ip6-address 2002:AC14:7865::
      end
   end
```

To configure system information on Router2 - web-based manager

1. Go to **System > Dashboard > Status**.
2. For **Host name**, select **Change**.
3. Enter “Router2”.
4. Go to **System > Config > Features**.
5. In **Basic Features**, enable **IPv6**, and select **Apply**.
6. Go to **System > Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	2002:A0B:6566::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	2002:AC14:7866::/0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet
Administrative Status	Up

To configure system information on Router2 - CLI

```
config system global
   set hostname Router2
   set gui-ipv6 enable
end
config system interface
   edit port1
      set alias internal
      set allowaccess https ping ssh
      set description “Internal RnD network”
      config ipv6
         set ip6-address 2002:a0b:6566::/0
      end
   next
   edit port2
      set alias ISP
      set allowaccess https ping ssh
```

```
      set description "ISP and Internet"
      config ipv6
         set ip6-address 2002:AC14:7866::
      end
   end
```

Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include: the internal network and the ISP network. There is no redistribution and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

To configure RIPng on Router1 - CLI

```
config router ripng
   config interface
      edit port1
      next
      edit port2
      end
   config neighbor
      edit 1
         set interface port1
         set ipv6 2002:a0b:6566::/0
      next
      edit 2
         set interface port2
         set ipv6 2002:AC14:7805::/0
      end
```

To configure RIPng on Router2 - CLI

```
config router ripng
   config interface
      edit port1
      next
      edit port2
      end
   config neighbor
      edit 1
         set interface port1
         set ipv6 2002:a0b:6565::/0
      next
      edit 2
         set interface port2
         set ipv6 2002:AC14:7805::/0
      end
```

Configuring other network devices

The other devices on the internal network all support IPv6 and are running RIPng, where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information, such as IPv6 addresses.

Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the FortiOS Handbook Troubleshooting chapter.

For troubleshooting problems with RIP, see RIPng: RIP and IPv6 on page 151.

Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems:

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit:

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table:

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous `diagnose ipv6 route list` command, but it is presented in a format that is easier to read.

```
get router info6 rip interface external
```

View the brief output on the RIP information for the interface listed. This includes information such as, if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

OSPF

This section describes Open Shortest Path First (OSPF) routing.

OSPF background and concepts

Open Shortest Path First (OSPF) is a link-state interior routing protocol that is widely used in large enterprise organizations. It only routes packets within a single autonomous system (AS). This is different from BGP, because BGP can communicate between ASs

Background

OSPF version 2 (OSPFv2) was defined in 1998 in RFC 2328. OSPF was designed to support classless IP addressing and variable subnet masks. This was a shortcoming of the earlier RIP protocols.

Updates to OSPFv2 are included in OSPF version 3 (OSPFv3), defined in 2008 in RFC 5340. OSPFv3 includes support for IPv6 addressing, where OSPF2 only supports IPv4 addressing.

The main benefit of OSPF is that it detects link failures in the network quickly and within seconds, has converged network traffic successfully without any networking loops. Also, OSPF has many features to control which routes are propagated and which are not, maintaining smaller routing tables. OSPF can also provide better load-balancing on external links than other interior routing protocols.

The parts and terminology of OSPF

The parts and terminology of OSPF include the following sections.

OSPFv3 and IPv6

OSPF version 3 (OSPFv3) includes support for IPv6. Generally, all IP addresses are in IPv6 format instead of IPv4. However, OSPFv3 area numbers use the same 32-bit numbering system as OSPFv2, as described in RFC 2740. Likewise, the router ID and area ID are in the same format as OSPFv2.

As with most advanced routing features on your FortiGate unit, IPv6 settings for dynamic routing protocols must be enabled before they are visible in the GUI. To enable IPv6 configuration in the GUI, enable it in **System > Feature Visibility**.

For IPv6, the main difference in OSPFv3 is that rather than using a network statement to enable OSPFv3 on an interface, you define OSPF6 (OSPF for IPv6) interfaces, which are bound to the interface and area. This configuration must be done in the CLI, as follows (with sample interfaces and addresses):

```
config router ospf6
   config area
      edit 0.0.0.0
   next
end
config ospf6-interface
   edit "tunnel"
   set interface "to_FGT300A-7"
```

```
next
   edit "internal_lan"
   set interface "port1"
next
   set router-id 10.174.0.113
end
```

Note that OSPFv3 neighbors use link-local IPv6 addresses, but with broadcast and point-to-point network types, and neighbors are automatically discovered. You only have to manually configure neighbors when using non-broadcast network types.

Router ID

In OSPF, each router has a unique 32-bit number that is called its router ID. Often, this 32-bit number is written the same as a 32-bit IPv4 address would be written in dotted decimal notation. However, some brands of routers, such as Cisco routers, support a router ID entered as an integer instead of an IP address.

It is a good idea not to use an IP address for the router ID that is already in use on the router. The router ID does not have to be a particular IP address on the router. By choosing a different number, it will be harder to get confused about which number you are looking at. It is a good idea to use as many of the area's numbers as possible. For example, if you have 15 routers in area 0.0.0.0, they could be numbered from 0.0.0.1 to 0.0.0.15. If you have an area 1.1.1.1, then routers in that area could start at 1.1.1.10.

You can manually set the router ID on your FortiGate unit:

To manually set an OSPF router ID of 0.0.1.1 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. For **Router ID**, enter `0.0.1.1`.
3. Select **Apply**.

To manually set an OSPF router ID of 0.0.1.1 - CLI

```
config router ospf
   set router-id 0.0.1.1
end
```

Adjacency

In an OSPF routing network, an OSPF router sends out OSPF hello packets when it boots up, to try to find any neighbours (routers that have access to the same network as the router booting up). Once neighbors are discovered and Hello packets are exchanged, updates are sent and the link state databases of both neighbors are synchronized. At this point, these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met:

- The subnet mask used on both routers must be the same subnet.
- The subnet number derived using the subnet mask and each router's interface IP address must match.
- The hello interval and the dead interval must match.
- The routers must have the same OSPF area ID. If they are in different areas, they are not neighbors.
- If authentication is used, they must pass authentication checks.

If any of these parameters are different between the two routers, the routers do not become OSPF neighbors and cannot be adjacent. If the routers become neighbors, they are adjacent.

Adjacency and neighbors

Neighbor routers can be in a two-way state, and not be adjacent. Adjacent routers normally have a neighbor state of FULL. Neighbors only exchange hello packets and do not exchange routing updates. Adjacent routers exchange LSAs (LSDB information) as well as hello packets. A good example of an adjacent pair of routers is the designated router (DR) and backup designated router (BDR).

You can check on the state of an OSPF neighbor using the CLI `get router info ospf neighbor all` command. For for more information, see OSPF background and concepts on page 158.

Why adjacency is important

It is important to have adjacent pairs of routers in the OSPF routing domain because routing protocol packets are only passed between adjacent routers. This means adjacency is required for two OSPF routers to exchange routes.

If there is no adjacency between two routers, such as one on the 172.20.120.0 network and another on the 10.11.101.0 network, the routers do not exchange routes. This makes sense because if all OSPF routers on the OSPF domain exchanged updates, it would flood the network.

Also, it is better for updates to progress through adjacent routers to ensure there are no outages along the way. Otherwise, updates could skip over routers that are potentially offline, causing longer routing outages and delays, while the OSPF domain learns of this outage later on.

If the OSPF network has multiple border routers and multiple connections to external networks, the designated router (DR) determines which router pairs become adjacent. The DR can accomplish this because it maintains the complete topology of the OSPF domain, including which router pairs are adjacent.

The backup designated router (BDR) also has this information in case the DR goes offline.

Designated router and backup router

In OSPF, a router can have a number of different roles to play.

A designated router (DR) is the designated broadcasting router interface for an AS. It looks after all of the initial contact and other routing administration traffic. Having only one router do all of this this greatly reduces the network traffic and collisions.

If something happens and the designated router goes offline, the backup designated router (BDR) takes over. An OSPF FortiGate unit interface can become either a DR or BDR. Both the DR and the BDR cover the same area, and are elected at the same time. The election process does not have many rules, but the exceptions can become complex.

Benefits

The OSPF concept of the designated router is a big step above RIP. With all RIP routers doing their own updates all the time, RIP suffers from frequent and sometimes unnecessary updates that can slow down your network. With OSPF, not only do routing changes only happen when a link state changes instead of any tiny change to the routing table, but the designated router reduces this overhead traffic even more.

However, smaller network topologies may have only a couple of routers besides the designated router. This may seem excessive, but it maintains the proper OSPF form and it will still reduce the administration traffic, but to a lesser extent than on a large network. Also, your network topology will be ready whenever you choose to expand your network.

DR and BDR election

An election chooses DR and BDR from all the available routers. The election is primarily based on the priority setting of the routers, where the highest priority becomes the DR and the second highest becomes the BDR. To resolve any ties, the router with the highest router ID wins. For example, a router with a router ID of 192.168.0.1 would win over a router with a router ID of 10.1.1.2.

The router priority can vary from 0 to 255, but at 0 a router will never become a DR or BDR. If a router with a higher priority comes online after the election, it must wait until the DR and BDR go offline before it becomes the DR.

If the original DR goes offline, but is then available when the BDR goes offline later on, the original DR will be promoted back to DR without an election leaving the new BDR as it is.

With your FortiGate unit, to configure the port1 interface to be a potential OSPF DR or BDR called `ospf_DR` on the network, you need to raise the priority of the router to a very high number, such as 250 out of 255. This will ensure the interface has a chance to be a DR, but will not guarantee that it will be one. To help ensure it becomes a DR, you should give the interface a low numbered IP address, such as 10.1.1.1 instead of 192.168.1.1 (but that is not part of this example). Enter the following command:

```
config router ospf
   config ospf-interface
   edit "ospf_DR"
      set priority 250
   end
end
```

Area

An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it is divided into these smaller areas for manageability.

If there are two or more routers that are viable within an area, there will always be a designated router (DR) and a backup designated router (BDR). For more information about these router roles, see Designated router and backup router on page 160.

Defining a private OSPF area involves the following:

- Assigning a 32-bit number to the area that is unique on your network
- Defining the characteristics of one or more OSPF areas
- Creating associations between the OSPF areas that you defined and the local networks to include in the OSPF area
- Adjusting the settings of OSPF-enabled interfaces, if required

IPv6 OSPF area numbers use the same 32-bit number notation as IPv4 OSPF.

If you are using the web-based manager to perform these tasks, follow the procedures summarized below.

FortiGate units support the four main types of OSPF areas:

- Backbone area
- Stub area

- NSSA
- Regular area

Backbone area

Every OSPF network has at least one AS, and every OSPF network has a backbone area. The backbone is the main area, and possibly the only area. All other OSPF areas are connected to a backbone area. This means if two areas want to pass routing information back and forth, that routing information will go through the backbone on its way between those areas. For this reason, the backbone not only has to connect to all other areas in the network, but also has to be uninterrupted in order to be able to pass traffic to all points of the network.

The backbone area is referred to as area 0 because it has an IP address of 0.0.0.0.

Stub area

A stub area is an OSPF area that receives no outside routes advertised into it. All routing in it is based on a default route. This essentially isolates it from outside areas.

Stub areas are useful for small networks that are part of a larger organization, especially if the networking equipment cannot handle routing large amounts of traffic passing through, or if there are other reasons to prevent outside traffic, such as security. For example, most organizations do not want their accounting department to be the center of their network with everyone's traffic passing through there. It would increase the security risks, slow down their network, and it generally does not make sense.

A variation on the stub area is the totally stubby area. It is a stub area that does not allow summarized routes.

NSSA

A not-so-stubby-area (NSSA) is a stub area that allows for external routes to be injected into it. While it still does not allow routes from external areas, it is not limited to using only the default route for internal routing.

Regular area

A regular area is what all the other ASs are, all the non-backbone, non-stub, and non-NSSA areas. A regular area generally has a connection to the backbone, does receive advertisements of outside routes, and does not have an area number of 0.0.0.0.

Authentication

In the OSPF packet header, there are two authentication-related fields: AuType and Authentication.

All OSPF packet traffic is authenticated. Multiple types of authentication are supported in OSPFv2. However in OSPFv3, there is no authentication built-in but it is assumed that IPsec will be used for authentication instead.

Packets that fail authentication are discarded.

Null authentication

Null authentication indicates there is no authentication being used. In this case, the 16-byte authentication field is not checked, and can be any value. However, checksumming is still used to locate errors. On your FortiGate, this is the `none` option for authentication.

Simple password authentication

Simple password refers to a standard plain text string of characters. The same password is used for all transactions on a network. The main use for this type of authentication is to prevent routers from accidently joining the network. Simple password authentication is vulnerable to many forms of attack, and is not recommended as a secure form of authentication.

Cryptographic authentication

Cryptographic authentication involves the use of a shared secret key to authenticate all router traffic on a network. The key is never sent over the network in the clear. A packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

Your FortiGate unit supports all three levels of authentication through the authentication keyword associated with creating an OSPF interface .

For example, to create an OSPF interface called `Accounting` on the port1 interface that is a broadcast interface, has a hello interval of 10 seconds, has a dead interval of 40 seconds, uses text authentication (simple password) with a password of “ospf_test”, enter the command:

```
config router ospf
   config ospf-interface
   edit Accounting
      set interface port1
      set network-type broadcast
      set hello-interval 10
      set dead-interval 40
      set authentication text
      set authentication-key "ospf_test"
   end
end
```

Hello and dead intervals

The OSPF Hello protocol is used to discover and maintain communications with neighboring routers.

Hello packets are sent out at a regular interval for this purpose. The DR sends out the hello packets. In a broadcast network, the multicast address of 224.0.0.5 is used to send out hello packets. New routers on the network listen for and reply to these packets to join the OSPF area. If a new router never receives a hello packet, other routers will not know it is there and will not communicate with it. However, once a new router is discovered, the DR adds it to the list of routers in that area and it is integrated into the routing calculations.

Dead interval is the time other routers will wait before declaring a neighbor dead (offline). It is very important to set a reasonable dead interval. If this interval is too short, routers will be declared offline when they are just slow or momentarily inaccessible, and link state updates will happen more than they need to, using more bandwidth. If the dead interval is too long, it will slow down network traffic overall if online routers attempt to contact offline ones instead of re-routing traffic.

FortiOS also supports OSPF fast-hello, which provides a way to send multiple hello packets per second. This is achieved by setting a dead-interval to one second. The hello-multiplier, which can be any number between 4 and 10, determines the number of hello packets that will be sent every second. The CLI syntax for OSPF fast-hello is as follows:

```
config ospf-interface
   edit ospf1
      set interface port1
      set network-type broadcast
      set dead-interval 1
      set hello-multiplier 4
end
```

Access lists

Access lists are filters used by FortiGate unit OSPF routing. An access list provides a list of IP addresses and the action to take for them. An access list essentially makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example, if you want all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate unit attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found, the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. It also eases troubleshooting because if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the OSPF+ IPv6 protocols, you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of `10.10.10.10` and `11.11.11.11`, enter the command:

```
config router access-list
   edit test_list
   config rule
      edit 1
         set prefix 10.10.10.10 255.255.255.255
         set action allow
         set exact-match enable
      next
      edit 2
         set prefix 11.11.11.11 255.255.255.255
         set action allow
         set exact-match enable
      end
   end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the access-list6 command as follows:

```
config router access-list6
   edit test_list_ip6
      config rule
         edit 1
            set prefix6 2002:A0A:A0A:0:0:0:0:0:/48
            set action deny
```

```
        next
        edit 2
           set prefix6 2002:B0B:B0B:0:0:0:0:0/48
           set action deny
        end
```

To use an access_list, you must call it from a routing protocol such as RIP. The following example uses the access_list from the earlier example called test_list to match routes coming in on the port1 interface. When there is a match, it will add 3 to the hop count metric for those routes to artificially decrease their priority. Enter the following command:

```
config router ospf
   config distribute-list
      edit 5
         set access-list test_list
         set protocol connected
      end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route 0.0.0.0/0 cannot be exactly matched with an access-list. A prefix-list must be used for this purpose.

How OSPF works

An OSPF installation consists of one or more areas. An OSPF area is typically divided into logical areas linked by Area Border Routers (ABR). A group of contiguous networks form an area. An ABR links one or more areas to the OSPF network backbone (area ID 0). For more information, see Dynamic routing overview on page 104.

OSPF is an interior routing protocol. It includes a backbone AS and possibly additional ASs. The DR and BDR are elected from potential routers with the highest priorities. The DR handles much of the administration to lower the network traffic required. New routers are discovered through hello packets sent from the DR using the multicast address of 224.0.0.5. If the DR goes offline at any time, the BDR has a complete table of routes that it uses when it takes over as the DR router.

OSPF does not use UDP or TCP, but is encapsulated directly in IP datagrams as protocol 89. This is in contrast to RIP and BGP. OSPF handles its own error detection and correction functions.

The OSPF protocol, when running on IPv4, can operate securely between routers, optionally using a variety of authentication methods to allow only trusted routers to participate in routing. OSPFv3, running on IPv6, no longer supports protocol-internal authentication. Instead, it relies on IPv6 protocol security (IPsec).

Other important parts of how OSPF works include:

- OSPF router discovery
- How OSPF works on FortiGate units
- External routes
- Link state database and route updates
- OSPF packets

OSPF router discovery

OSPF-enabled routers generate link state advertisements (LSA) and send them to their neighbors whenever the status of a neighbor changes or a new neighbor comes online. As long as the OSPF network is stable, LSAs between OSPF neighbors do not occur. An LSA identifies the interfaces of all OSPF-enabled routers in an area,

and provides information that enables OSPF-enabled routers to select the shortest path to a destination. All LSA exchanges between OSPF-enabled routers are authenticated.

When a network of OSPF routers comes online, the following steps occur:

1. When OSPF routers come online, they send out hello packets to find other OSPF routers on their network segment.
2. When they discover other routers on their network segment, they generally become adjacent. Adjacent routers can exchange routing updates. For more information, see Adjacency on page 159.
3. A DR and BDR are elected from the available routers using priority settings and router ID. See Designated router and backup router on page 160, and OSPF background and concepts on page 158.
4. Link state updates are sent between adjacent routers to map the topology of the OSPF area.
5. Once complete, the DR floods the network with the updates to ensure all OSPF routers in the area have the same OSPF route database. After the initial update, there are very few required updates if the network is stable.

How OSPF works on FortiGate units

When a FortiGate unit interface is connected to an OSPF area, that unit can participate in OSPF communications. FortiGate units use the OSPF hello protocol to acquire neighbors in an area. A neighbor is any router that is directly connected to the same area as the FortiGate unit and ideally is adjacent with a state of Full. After initial contact, the FortiGate unit exchanges hello packets with its OSPF neighbors regularly to confirm that the neighbors can be reached.

The number of routes that a FortiGate unit can learn through OSPF depends on the network topology. A single unit can support tens of thousands of routes if the OSPF network is configured properly.

External routes

OSPF is an internal routing protocol. OSPF external routes are routes where the destination is using a routing protocol other than OSPF. OSPF handles external routes by adjusting the cost of the route to include the cost of the other routing protocol. There are two methods of calculating this cost, which are used for OSPF external1 (E1) and OSPF external2 (E2).

OSPF E1

In OSPF E1, the destination is outside the OSPF domain. This requires a different metric to be used beyond the normal OSPF metrics. The new metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.

OSPF E2

OSPF E2 is the default external type when routes are redistributed outside of OSPF. With OSPF E2, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. Dropping the OSPF portion can be useful in a number of situations, for example, on border routers that have no OSPF portion or where the OSPF routing cost is negligible compared to the external routing cost.

Comparing E1 and E2

The best way to understand OSPF E1 and E2 routes is to check routing tables on OSPF routers. If you look at the routes on an OSPF border router, the redistributed routes will have an associated cost that represents only the external route, as there is no OSPF cost to the route due to it already being on the edge of the OSPF domain. However, if you look at that same route on a different OSPF router inside the OSPF routing domain, it will have a

higher associated cost, essentially the external cost plus the cost over the OSPF domain to that border router. The border router uses OSPF E2, where the internal OSPF router uses OSPF E1 for the same route.

Viewing external routes

When you are trying to determine the costs for routes in your network to predict how traffic will be routed, you need to see the external OSPF routes and their associated costs. On your FortiGate unit, you find this information through the CLI.

To view external routes - CLI

You can view the whole routing table using `get router info routing-table all` to see all the routes, including the OSPF external routes. For a shorter list, you can use the `get router info routing-table ospf` command. The letter at the left will be either E1 or E2 for external OSPF routes. The output will look similar to the following, depending on what routes are in your routing table:

```
FGT620B# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

O*E2   0.0.0.0/0 [110/10] via 10.1.1.3, tunnel_wan2, 00:02:11
O      10.0.0.1/32 [110/300] via 10.1.1.3, tunnel_wan2, 00:02:11
S      0.0.0.0/0 [10/0] via 192.168.183.254, port2
S      1.0.0.0/8 [10/0] via 192.168.183.254, port2
```

Link state database and route updates

OSPF is based on links. The links between adjacent neighbor routers allow updates to be passed along the network. Network links allow the DR to flood the area with link state database (LSDB) updates. External links allow the OSPF area to connect to destinations outside the OSPF autonomous system. Information about these links is passed throughout the OSPF network as link state updates.

The LSDB contains the information that defines the complete OSPF area, but the LSDB is not the routing table. It contains the information from all the link state updates passed along the network. When there are no more changes required and the network is stable, the LSDB on each router in the network will be the same. The DR will flood the LSDB to the area to ensure each router has the same LSDB.

To calculate the best route (shortest path) to a destination, the FortiGate unit applies the Shortest Path First (SPF) algorithm, based on Dijkstra's algorithm, to the accumulated link state information. OSPF uses relative path cost metric for choosing the best route. The path cost can be any metric, but is typically the bandwidth of the path, which is how fast traffic will get from one point to another.

The path cost, similar to "distance" for RIP, imposes a penalty on the outgoing direction of a FortiGate unit interface. The path cost of a route is calculated by adding all of the costs associated with the outgoing interfaces along the path to the destination. The lowest overall path cost indicates the best route, and generally the fastest route. Some brands of OSPF routers, such as Cisco, implement cost as a direct result of bandwidth between the routers. Generally this is a good cost metric because larger bandwidth means more traffic can travel without slowing down. To achieve this type of cost metric on FortiGate units, you need to set the cost for each interface manually in the CLI.

The inter-area routes may not be calculated when a Cisco type ABR has no fully adjacent neighbor in the backbone area. In this situation, the router considers summary-LSAs from all Actively summary-LSAs from all Actively Attached areas (RFC 3509).

The FortiGate unit dynamically updates its routing table based on the results of the SPF calculation to ensure that an OSPF packet will be routed using the shortest path to its destination. Depending on the network topology, the entries in the FortiGate unit routing table may include:

- The addresses of networks in the local OSPF area (to which packets are sent directly)
- Routes to OSPF area border routers (to which packets destined for another area are sent)
- Routes to area boundary routers, if the network contains OSPF areas and non-OSPF domains, which reside on the OSPF network backbone and are configured to forward packets to destinations outside the OSPF AS.

OSPF route updates

Once the OSPF domain is established, there should be few updates required on a stable network. When updates occur and a decision is required concerning a new route, this is the general procedure.

Our router gets a new route and needs to decide if it should go in the routing table.

The router has an up-to-date LSDB of the entire area, containing information about each router, the next hop to it, and most importantly the cost to get there.

Our router turns the LSDB into an SPF tree using Dijkstra's algorithm. It does not matter if there is more than one path to a router on the network, the SPF tree only cares about the shortest path to that router.

Once the SPF tree has been created and shows the shortest paths to all the OSPF routers on the network, the work is done. If the new route is the best route, it will be part of that tree. If it is not the shortest route, it will not be included in the LSDB.

If there has been a change from the initial LSDB to the new SPF tree, a link state update will be sent out to let the other routers know about the change so they can also update their LSDBs. This is vital since all routers on the OSPF area must have the same LSDB.

If there was no change between the LSDB and the SPF tree, no action is taken.

OSPF packets

Every OSPF packet starts with a standard 24-byte header, and another 24 bytes of information or more. The header contains all the information necessary to determine whether the packet should be accepted for further processing.

OSPF packet

1-byte Version field	1-byte Type field	2-byte Packet length	3-byte Router ID
4-byte Area ID	2-byte Checksum	2-byte Auth Type	8-byte Authentication
4-byte Network Mask	2-byte Hello interval	1-byte Options field	1-byte Router Priority
4-byte Dead Router interval	4-byte DR field	4-byte BDR field	4-byte Neighbor ID

The following descriptions summarize the OSPF packet header fields:

Version field: The OSPF version number. This specification documents version 2 of the protocol.

Type field: There are 5 OSPF packet types. From one to five, respectively, they are Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment.

Packet length: The length of the OSPF protocol packet, in bytes. This length includes the standard OSPF 24-byte header, so all OSPF packets are at 24-bytes long.

Router ID: The Router ID of the packet's source.

Area ID: A 32-bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.

Checksum: The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field. If the packet's length is not an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. The checksum is considered to be part of the packet authentication procedure. For some authentication types, the checksum calculation is omitted.

Auth Type: Identifies the authentication procedure to be used for the packet. Authentication types include Null authentication (0), Simple password (1), Cryptographic authentication (2), and all others are reserved for future use.

Authentication: A 64-bit field for use by the authentication scheme. When AuType indicates no authentication is being used, the authentication field is not checked and can be any value. When AuType is set to 2 (cryptographic authentication), the 64-bit authentication field is split into the following four fields: Zero field, Key ID field, Authentication data length field, and Cryptographic sequence field.

The Key ID field indicates the key and algorithm used to create the message digest appended to the packet. The Authentication data length field indicates how many bytes long the message digest is. The Cryptographic sequence field is a non-decreasing number that is set when the packet is received and authenticated to prevent replay attacks.

Network Mask: The subnet where this packet is valid.

Hello interval: The period of time between sending out hello packets. For more information, see Hello and dead intervals on page 163.

Options field: The OSPF protocol defines several optional capabilities. A router indicates the optional capabilities that it supports in its OSPF hello packets, database description packets and in its LSAs. This enables routers supporting a mix of optional capabilities to coexist in a single AS.

Router priority: The priority, between 0 and 255, that determines which routers become the DR and BDR. For more information, see Designated router and backup router on page 160.

Dead router interval: The period of time when there is no response from a router before it is declared dead. For more information, see Hello and dead intervals on page 163.

DR and BDR fields: The DR and BDR fields each list the router that fills that role on this network, generally the routers with the highest priorities. For more information, see Designated router and backup router on page 160.

Neighbor ID: The ID number of a neighboring router. This ID is used to discover new routers and respond to them.

Troubleshooting OSPF

As with other dynamic routing protocols, OSPF has some issues that may need troubleshooting from time to time. For basic troubleshooting, see the FortiOS Handbook Troubleshooting chapter.

Clearing OSPF routes from the routing table

If you think the wrong route has been added to your routing table and you want to check it out, you first have to remove that route from your table before seeing if it is added back in or not. You can clear all or some OSPF neighbor connections (sessions) using the `execute router clear ospf` command. The exec router clear command is much more limiting for OSPF than it is for BGP. For more information, see BGP on page 199.

For example, if you have routes in the OSPF routing table and you want to clear the specific route to IP address 10.10.10.1, you will have to clear all the OSPF entries. Enter the command:

```
execute router clear ospf process
```

Checking the state of OSPF neighbors

In OSPF, each router sends out link state advertisements to find other routers on its network segment and to create adjacencies with some of those routers. This is important because routing updates are only passed between adjacent routers. If two routers you believe to be adjacent are not, that can be the source of routing failures.

To identify this problem, you need to check the state of the OSPF neighbors of your FortiGate unit. Use the CLI `get router info ospf neighbor all` command to see all the neighbors for your FortiGate unit. You will see output in the form of:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID  Pri  State    Dead Time Address Interface
10.0.0.2     1    Full/ -  00:00:39 10.1.1.2 tunnel_wan1
10.0.0.2     1    Full/ -  00:00:34 10.1.1.4 tunnel_wan2
```

The important information here is the `State` column. Any neighbors that are not adjacent to your FortiGate unit will be reported in this column as something other than `Full`. If the state is `Down`, that router is offline.

Passive interface problems

A passive OSPF interface does not send out any updates. This means it cannot be a DR, BDR, or an area border router among other things. It will depend on other neighbor routers to update its link state table.

Passive interfaces can cause problems when they are not receiving the routing updates you expect from their neighbors. This will result in the passive OSPF FortiGate unit interface having an incomplete or out-of-date link state database, and it will not be able to properly route its traffic. It is possible that the passive interface is causing a hole in the network where no routers are passing updates to each other, however, this is a rare situation.

If a passive interface is causing problems, there are simple methods to determine it is the cause. The easiest method is to make it an active interface, and if the issues disappear, then that was the cause. Another method is

to examine the OSPF routing table and related information to see if it is incomplete compared to other neighbor routers. If this is the case, you can clear the routing table, reset the device, and allow it to repopulate the table.

If you cannot make the interface active for some reason, you will have to change your network to fix the “hole” by adding more routers, or changing the relationship between the passive router’s neighbors to provide better coverage.

Timer problems

A timer mismatch is when two routers have different values set for the same timer. For example, if one router declares a router dead after 45 seconds and another waits for 4 minutes, that difference in time will result in those two routers being out of synch for that period of time. One will still see the offline router as being online.

The easiest method to check the timers is to check the configuration on each router. Another method is to sniff some packets, and read the timer values in the packets themselves from different routers. Each packet contains the hello interval and dead interval periods, so you can compare them easily enough.

BFD

Bidirectional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

Authentication issues

OSPF has a number of authentication methods you can choose from. You may encounter problems with routers not authenticating as you expect. This will likely appear simply as one or more routers that have a blind spot in their routing and they will not acknowledge a router. This can be a problem if that router connects areas to the backbone, as it will appear to be offline and unusable.

To confirm this is the issue, the easiest method is to turn off authentication on the neighboring routers. With no authentication between any routers, everything should flow normally.

Another method to confirm that authentication is the problem is to sniff packets and look at their contents. The authentication type and password are right in the packets which makes it easy to confirm they are what you expect during real time. It is possible one or more routers is not configured as you expect and may be using the wrong authentication. This method is especially useful if there are a group of routers with these problems since it may be only one router causing the problem that is seen in multiple routers.

Once you have confirmed the problem is related to authentication, you can decide how to handle it. You can turn off authentication and take your time to determine how to get your preferred authentication type back online. You can try another type of authentication, such as text instead of md5, which may have more success and still provide some level of protection. The important part is that once you confirm the problem, you can decide how to fix it properly.

DR and BDR election issues

You can force a particular router to become the DR and BDR by setting its priorities higher than any other OSPF routers in the area. This is a good idea when those routers have more resources to handle the traffic and extra work of the DR and BDR roles, since not all routers may be able to handle all of that traffic.

However, if you set all the other routers so they do not have a chance at being elected (give them a priority of 0), you can run into problems if the DR and BDR go offline. The good part is that you will have some warning generally as the DR goes offline and the BDR is promoted to the DR position. However, if the network segment with both the DR and BDR goes down, your network will have no way to send hello packets, send updates, or perform the other tasks that the DR performs.

The solution to this is to always allow routers to have a chance to be promoted, even if you set their priority to 1. In that case, they will be the last choice but if there are no other candidates, you want that router to become the DR. Most networks will have already alerted you to the equipment problems, so this will be a temporary measure to keep the network traffic moving until you can find and fix the problem and get the real DR back online.

Basic OSPF example

This example sets up an OSPF network at a small office. There are 3 routers, all running OSPFv2. The border router connects to a BGP network.

All three routers in this example are FortiGate units. Router1 will be the designated router (DR) and Router2 will be the backup designated router (BDR) due to their priorities. Router3 will not be considered for either the DR or BDR elections. Instead, Router3 is the Autonomous System Border Router (ASBR) routing all traffic to the ISP's BGP router on its way to the Internet.

Router2 has a modem connected that provides dialup access to the Internet as well, at a reduced bandwidth. This is a PPPoE connection to a DSL modem. This provides an alternate route to the Internet if the other route goes down. The DSL connection is slow and is charged by the amount of traffic. For these reasons, OSPF will highly favor Router3's Internet access.

The DSL connection connects to an OSPF network with the ISP, so no redistribution of routes is required. However, the ISP network does have to be added to that router's configuration.

Network layout and assumptions

There are three FortiGate units acting as OSPFv2 routers on the network: Router1, Router2, and Router3. Router1 will be the DR, and Router 2 the BDR. Router3 is the ASBR that connects to the external ISP router running BGP. Router2 has a PPPoE DSL connection that can access the Internet.

The head office network is connected to Router1 and Router2 on the 10.11.101.0 subnet.

Router1 and Router3 are connected over the 10.11.103.0 subnet.

Router2 and Router3 are connected over the 10.11.102.0 subnet.

The following table lists the router, interface, address, and role it is assigned.

Routers, interfaces, and IP addresses for the basic OSPF example network

Router name	Interface	IP address	Interface is connected to:
Router1 (DR)	Internal (port1)	10.11.101.1	Head office network and Router2
	External (port2)	10.11.102.1	Router3

Router name	Interface	IP address	Interface is connected to:
Router2 (BDR)	Internal (port1)	10.11.101.2	Head office network and Router1
	External (port2)	10.11.103.2	Router3
	DSL (port3)	10.12.101.2	PPPoE DSL access
	Internal1 (port1)	10.11.102.3	Router1
Router3 (ASBR)	Internal2 (port2)	10.11.103.3	Router2
	External (port3)	172.20.120.3	ISP's BGP network

Basic OSPF network topology

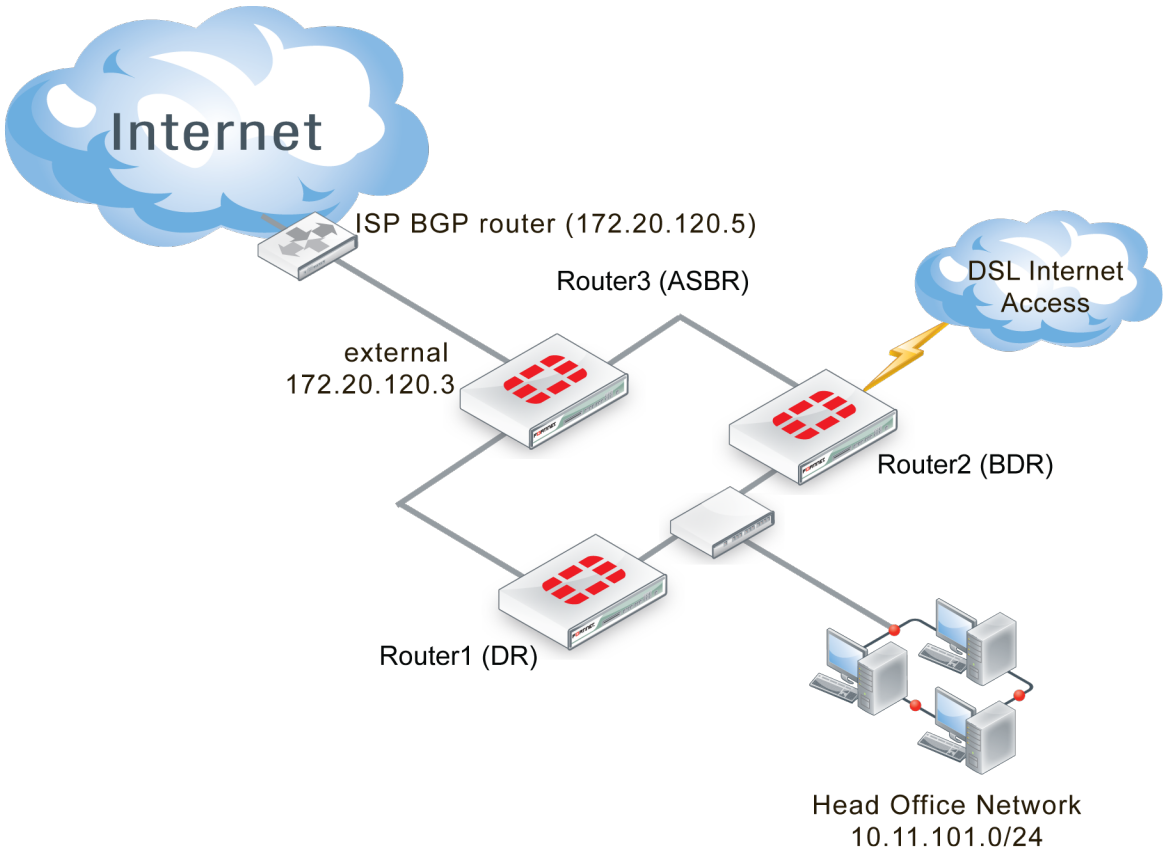

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed and are in NAT/Route operation mode.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.

- This OSPF network is not connected to any other OSPF networks.
- Both Internet connections are always available.
- The modem connection is very slow and expensive.
- Other devices may be on the network, but do not affect this basic configuration.
- Router3 is responsible for redistributing all routes into and out of the OSPF AS.

Configuring the FortiGate units

Each FortiGate unit needs the interfaces and basic system information, such as hostname, configured.

This section includes:

- Configuring Router1
- Configuring Router2
- Configuring Router3

Configuring Router1

Router1 has two interfaces connected to the network: internal (port1) and external (port2). Its host name must be changed to Router1.

To configure Router1 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Beside the host name, select **Change**.
3. Enter a hostname of `Router1` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router2
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	External
IP/Network Mask	10.11.102.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

Configuring Router2

Router2 configuration is the same as Router1, except Router2 also has the DSL interface to configure.

The DSL interface is configured with a username of “user1” and a password of “ospf_example”. The default gateway will be retrieved from the ISP and the defaults will be used for the rest of the PPPoE settings.

To configure Router2 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Beside the host name, select **Change**.
3. Enter a hostname of `Router2` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router1
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	External
IP/Network Mask	10.11.103.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

6. Edit DSL (port3), set the following information and select **OK**.

Alias	DSL
Addressing Mode	PPPoE
Username	user1
Password	ospf_example
Unnumbered IP	10.12.101.2/255.255.255.0
Retrieve default gateway from server	Enable

Administrative Access	HTTPS SSH PING
Description	DSL
Administrative Status	Up

Configuring Router3

Router3 is similar to Router1 and Router2 configurations. The main difference is the External (port3) interface connected to the ISP BGP network, which has no administration access enabled for security reasons.

To configure Router3 interfaces - web-based manager

1. Go to **System > Status > Dashboard**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router3` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.102.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router1
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	Internal2
IP/Network Mask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2
Administrative Status	Up

6. Edit port3, set the following information and select **OK**.

Alias	External
IP/Network Mask	172.20.120.3/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	ISP BGP
Administrative Status	Up

Configuring OSPF on the FortiGate units

With the interfaces configured, now the FortiGate units can be configured for OSPF on those interfaces. All routers are part of the backbone 0.0.0.0 area, so there is no inter-area communications needed.

For a simple configuration, there will be no authentication, no graceful restart or other advanced features, and timers will be left at their defaults. Also, the costs for all interfaces will be left at 10, except for the modem and ISP interfaces where cost will be used to load balance traffic. Nearly all advanced features of OSPF are only available from the CLI.

The network that is defined covers all the subnets used in this example - 10.11.101.0, 10.11.102.0, and 10.11.103.0. All routes for these subnets will be advertised. If there are other interfaces on the FortiGate units that you do not want included in the OSPF routes, ensure those interfaces use a different subnet outside of the 10.11.0.0 network. If you want all interfaces to be advertised you can use an OSPF network of 0.0.0.0 .

Each router will configure:

- Router ID
- Area
- Network
- Two or three interfaces depending on the router
- Priority for DR (Router1) and BDR (Router2)
- Redistribute for ASBR (Router3)

This section includes:

- Configuring OSPF on Router1
- Configuring OSPF on Router2
- Configuring OSPF on Router3

Configuring OSPF on Router1

Router1 has a very high priority to ensure it becomes the DR for this area. Also Router1 has the lowest IP address to help ensure it will win in case there is a tie at some point. Otherwise, it is a standard OSPF configuration. Setting the priority can only be done in the CLI, and it is for a specific OSPF interface.

To configure OSPF on Router1 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.1` and select **Apply**.
3. In **Areas**, select **Create New,** set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router1-Internal-DR
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

6. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router1-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

7. Using the CLI, enter the following commands to set the priority for the Router1-Internal OSPF interface to maximum, ensuring this interface becomes the DR.

```
config router ospf
   config ospf-interface
      edit Router1-Internal-DR
         set priority 255
      end
```

To configure OSPF on Router1 - CLI

```
config router ospf
   set router-id 10.11.101.1
   config area
      edit 0.0.0.0
```

```
        next
    end
    config network
        edit 1
            set prefix 10.11.0.0/255.255.255.0
        next
    end
    config ospf-interface
        edit "Router1-Internal"
            set interface "port1"
            set priority 255
        next
        edit "Router1-External"
            set interface "port2"
        next
    end
end
```

Configuring OSPF on Router2

Router2 has a high priority to ensure it becomes the BDR for this area and configures the DSL interface slightly differently. Assume this will be a slower connection resulting in the need for longer timers and a higher cost for this route.

Otherwise, it is a standard OSPF configuration.

To configure OSPF on Router2 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.2` and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Router2-Internal
Interface	port1 (Internal)
IP	0.0.0.0

Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	

6. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router2-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

7. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Router2-DSL
Interface	port3 (DSL)
IP	0.0.0.0
Authentication	none
Cost	50
Timers (seconds)	
Hello Interval	20
Dead Interval	80

8. Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR:

```
config router ospf
   config ospf-interface
      edit Router2-Internal
         set priority 250
      next
   end
```

To configure OSPF on Router2 - CLI

```
config router ospf
   set router-id 10.11.101.2
   config area
      edit 0.0.0.0
      next
   end
   config network
      edit 1
         set prefix 10.11.0.0/255.255.0.0
      next
   end
   config ospf-interface
      edit "Router2-Internal"
         set interface "port1"
         set priority 255
      next
      edit "Router2-External"
         set interface "port2"
      next
      edit "Router2-DSL"
         set interface "port3"
         set cost 50
      next
   end
end
```

Configuring OSPF on Router3

Router3 is more complex than the other two routers. The interfaces are straightforward, but this router has to import and export routes between OSPF and BGP. That requirement makes Router3 an ASBR. Also, Router3 needs a lower cost on its route to encourage all traffic to the Internet to route through it.

In the advanced OSPF options, redistribute is enabled for Router3. It allows different types of routes, learned outside of OSPF, to be used in OSPF. Different metrics are assigned to these other types of routes to make them more or less preferred to regular OSPF routes.

To configure OSPF on Router3 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.2` and select **Apply**.
3. Expand **Advanced Options**.
4. In **Redistribute**, set the following information, and select **OK**.

Route type	Redistribute	Metric
Connected	Enable	15
Static	Enable	15
RIP	Disable	n/a
BGP	Enable	5

5. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router3-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router3-Internal2
Interface	port2 (Internal2)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router3-ISP-BGP

Interface	port3 (ISP-BGP)
IP	0.0.0.0
Authentication	none
Cost	2
Timers (seconds)	
Hello Interval	20
Dead Interval	80

10. Using the CLI, enter the following commands to set the priority for the Router3-Internal OSPF interface to ensure this interface will become the BDR.

```
config router ospf
   config ospf-interface
   edit Router3-Internal
      set priority 250
   next
end
```

To configure OSPF on Router3 - CLI

```
config router ospf
   set router-id 10.11.102.3
   config area
      edit 0.0.0.0
      next
   end
   config network
      edit 1
         set prefix 10.11.0.0/255.255.255.0
      next
      edit 2
         set prefix 172.20.120.0/255.255.255.0
      next
   end
   config ospf-interface
      edit "Router3-Internal"
         set interface "port1"
         set priority 255
      next
      edit "Router3-External"
         set interface "port2"
      next
      edit "Router3-ISP-BGP"
         set interface "port3"
         set cost 2
      next
   end
end
```

Configuring other networking devices

The other networking devices required in this configuration are on the two ISP networks, the BGP network for the main Internet connection, and the DSL backup connection.

In both cases, the ISPs need to be notified about the OSPF network settings including router IP addresses, timer settings, and so on. The ISP will use this information to configure its routers that connect to this OSPF network.

Testing network configuration

Testing the network configuration involves two parts: testing the network connectivity and testing the OSPF routing.

To test the network connectivity, use ping, traceroute, and other network tools.

To test the OSPF routing in this example, refer to the troubleshooting outlined in Basic OSPF example on page 172.

Advanced inter-area OSPF example

This example sets up an OSPF network at a large office. There are three areas, each with two routers. Typically OSPF areas would not be this small, and if they were, the areas would be combined into one larger area. However, the stub area services the accounting department whose members are very sensitive about their network and do not want their network information broadcasted through the rest of the company. The backbone area contains the bulk of the company's network devices. The regular area was established for various reasons, such as hosting the company servers in a separate area with extra security.

One area is a small stub area that has no independent Internet connection, and has only one connection to the backbone area. That connection between the stub area and the backbone area is only through a default route. No routes outside the stub area are advertised into that area. Another area is the backbone, which is connected to the other two areas. The third area has the Internet connection, and all traffic to and from the Internet must use that area’s connection. If that traffic comes from the stub area, then that traffic is treating the backbone like a transit area that only uses it to get to another area.

In the stub area, a subnet of computers is running the RIP routing protocol and those routes must be redistributed into the OSPF areas.

Network layout and assumptions

There are four FortiGate units in this network topology, which are acting as OSPF routers:

Advanced inter-area OSPF network topology

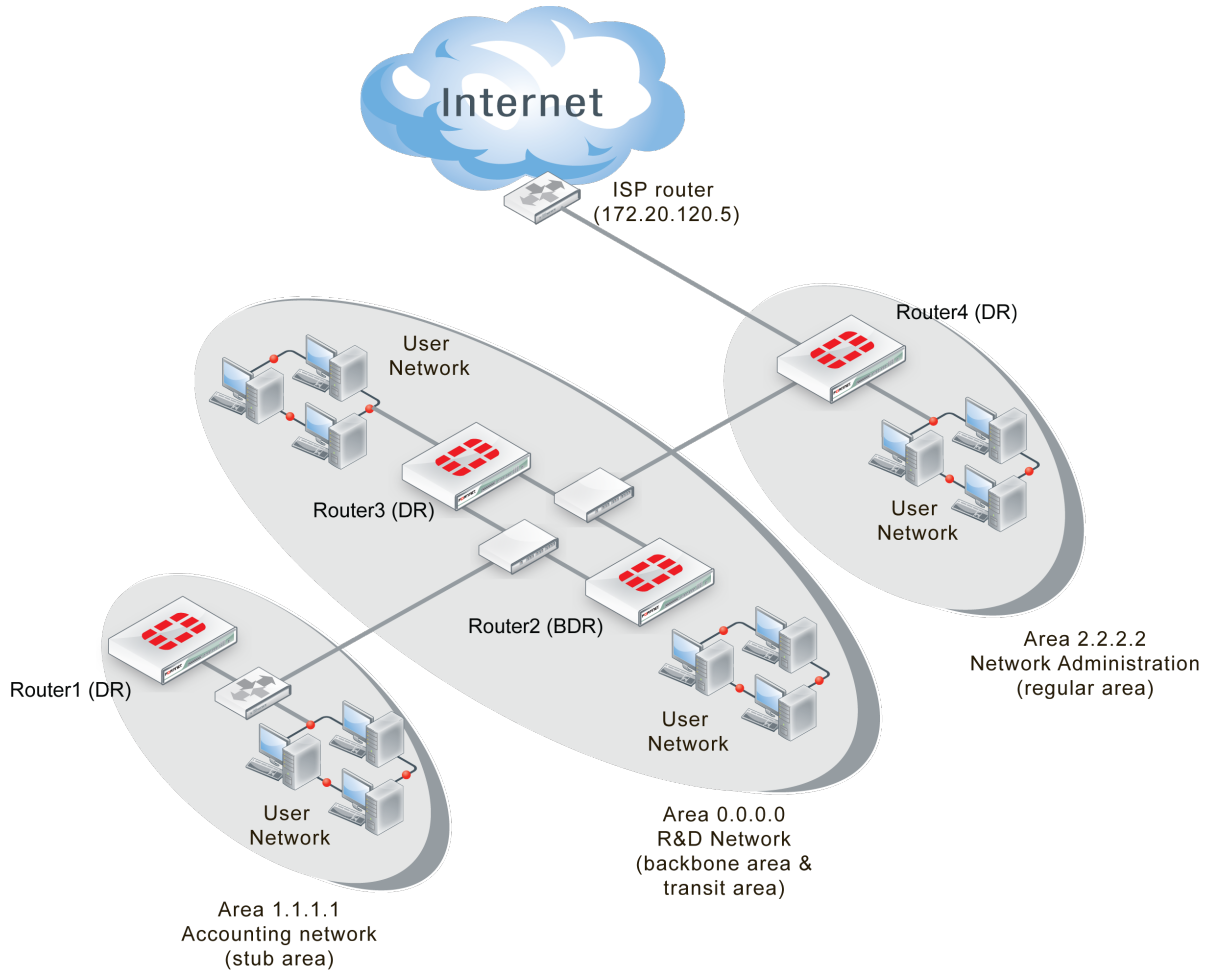

Area 1.1.1.1 is a stub area with one FortiGate unit OSPF router called Router1 (DR). Its only access outside of that area is a default route to the backbone area, which is how it accesses the Internet. Traffic must go from the stub area, through the backbone, to the third area to reach the Internet. The backbone area in this configuration is called a transit area. Also, in area 1.1.1.1 there is a RIP router that will be providing routes to the OSPF area through redistribution.

Area 0.0.0.0 is the backbone area and has two FortiGate unit routers named Router2 (BDR) and Router3 (DR).

Area 2.2.2.2 is a regular area that has an Internet connection accessed by both the other two OSPF areas. There is only one FortiGate unit router in this area called Router4 (DR). This area is more secure and requires MD5 authentication by routers.

All areas have user networks connected but they are not important for configuring the network layout for this example.

Internal interfaces are connected to internal user networks only. External1 interfaces are connected to the 10.11.110.0 network, joining Area 1.1.1.1 and Area 0.0.0.0.

External2 interfaces are connected to the 10.11.111.0 network, joining Area 0.0.0.0 and Area 2.2.2.2. The ISP interface is called ISP.

Routers, areas, interfaces, and IP addresses for advanced OSPF network

Router name	Area number and type	Interface	IP address
Router1 (DR)	1.1.1.1 - stub area (Accounting)	port1 (internal)	10.11.101.1
		port2 (external1)	10.11.110.1
Router2 (BDR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.102.2
		port2 (external1)	10.11.110.2
		port3 (external2)	10.11.111.2
Router3 (DR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.103.3
		port2 (external1)	10.11.110.3
		port3 (external2)	10.11.111.3
Router4 (DR)	2.2.2.2 - regular area (Network Admin)	port1 (internal)	10.11.104.4
		port2 (external2)	10.11.111.4
		port3 (ISP)	172.20.120.4

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate units used in this example have interfaces named port1, port2, and port3.
- All FortiGate units in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed and are in NAT/Route operation mode.
- During configuration, if settings are not directly referred to, they will be left at the default settings.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF areas outside of this example.
- The Internet connection is always available.
- Other devices may be on the network but do not affect this configuration.

Configuring the FortiGate units

This section configures the basic settings on the FortiGate units to be OSPF routers. These configurations include multiple interface settings and the hostname.

There are four FortiGate units in this example. The two units in the backbone area can be configured exactly the same except for IP addresses, so only the Router3 (the DR) configuration will be given, with notes indicating Router2's (the BDR) IP addresses.

Configuring the FortiGate units include:

- Configuring Router1
- Configuring Router2
- Configuring Router3
- Configuring Router4

Configuring Router1

Router1 is part of the Accounting network stub area (1.1.1.1).

To configure Router1 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router1` and select **OK**.
4. Go to **Network > Interfaces** edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	External1
IP/Network Mask	10.11.110.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone network and Internet
Administrative Status	Up

Configuring Router2

Router2 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

Router2 has three interfaces configured: one to the internal network and two to Router3 for redundancy.

To configure Router2 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router2` and select **OK**.
4. Go to **Network > Interfaces**, edit port1 (internal), set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.102.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information and select **OK**.

Alias	external1
IP/Network Mask	10.11.110.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.111.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 second connection
Administrative Status	Up

Configuring Router3

Router3 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

To configure Router3 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router3` and select **OK**.
4. Go to **Network > Interfaces**, edit port1 (internal), set the following information, and select **OK**.

Alias	internal

IP/Network Mask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information and select **OK**.

Alias	external1
IP/Network Mask	10.11.110.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.111.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 second connection
Administrative Status	Up

Configuring Router4

Router4 is part of the Network Administration regular area (2.2.2.2). This area provides Internet access for both area 1.1.1.1 and the backbone area.

This section configures interfaces and hostname.

To configure Router4 interfaces - web-based manager

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router4` and select **OK**.
4. Go to **Network > Interfaces**.
5. Edit port1 (internal).
6. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

7. Edit port2 (external2).
8. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.110.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone and Accounting network
Administrative Status	Up

9. Edit port3 (ISP).
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet
Administrative Status	Up

Configuring OSPF on the FortiGate units

Three of the routers are designated routers (DR) and one is a backup DR (BDR). This is achieved through the lowest router ID numbers, or OSPF priority settings.

Also, each area needs to be configured as each respective type of area: stub, backbone, or regular. This affects how routes are advertised into the area.

To configure OSPF on Router1 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Enter `10.11.101.1` for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New,** set the following information, and select **OK**.

Area	1.1.1.1
Type	Stub
Authentication	None

4. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.101.0/255.255.255.0
Area	1.1.1.1

5. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Accounting
Interface	port1 (internal)
IP	10.11.101.1
Authentication	None

6. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.1
Authentication	None

To configure OSPF on Router2 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Enter `10.11.102.2` for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New,** set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	None

4. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.102.2/255.255.255.0
Area	0.0.0.0

5. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.110.2/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.111.2/255.255.255.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	RnD network
Interface	port1 (internal)
IP	10.11.102.2
Authentication	None

8. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.2
Authentication	None

9. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.2
Authentication	None

To configure OSPF on Router3 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Enter `10.11.103.3` for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0

Type	Regular
Authentication	None

4. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.102.3/255.255.255.0
Area	0.0.0.0

5. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.110.3/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.111.3/255.255.255.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	RnD network
Interface	port1 (internal)
IP	10.11.103.3
Authentication	None

8. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.3
Authentication	None

9. In **Interfaces**, select **Create New,** set the following information, and select **OK**.

Name	Backbone2
Interface	port3 (external2)

IP	10.11.111.3
Authentication	None

To configure OSPF on Router4 - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. Enter `10.11.104.4` for the **Router ID** and then select **Apply**.
3. In **Areas**, select **Create New**.
4. Set the following information and select **OK**.

Area	2.2.2.2
Type	Regular
Authentication	None

5. In **Networks**, select **Create New,** set the following information, and select **OK**.

IP/Netmask	10.11.104.0/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.111.0/255.255.255.0
Area	0.0.0.0

7. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	172.20.120.0/255.255.255.0
Area	0.0.0.0

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Network Admin network
Interface	port1 (internal)
IP	10.11.104.4
Authentication	None

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone2

Interface	port2 (external2)
IP	10.11.111.4
Authentication	None

10. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	ISP
Interface	port3 (ISP)
IP	172.20.120.4
Authentication	None

Configuring other networking devices

All network devices on this network are running OSPF routing. The user networks (Accounting, R&D, and Network Administration) are part of one of the three areas.

The ISP needs to be notified of your network configuration for area 2.2.2.2. Your ISP will not advertise your areas externally as they are intended as internal areas. External areas have assigned unique numbers. The area numbers used in this example are similar to the 10.0.0.0 and 192.168.0.0 subnets used in internal networking.

Testing network configuration

There are two main areas to test in this network configuration: network connectivity and OSPF routing.

To test network connectivity, see if computers on the Accounting or R&D networks can access the Internet. If you need troubleshooting network connectivity, see the FortiOS Handbook Troubleshooting chapter.

To test OSPF routing, check the routing tables on the FortiGate units to ensure the expected OSPF routes are present. If you need help troubleshooting OSPF routing, see Advanced inter-area OSPF example on page 184.

Controlling redundant links by cost

In this scenario, two FortiGate units have redundant links: one link between their WAN1 interfaces and another between their WAN2 interfaces.

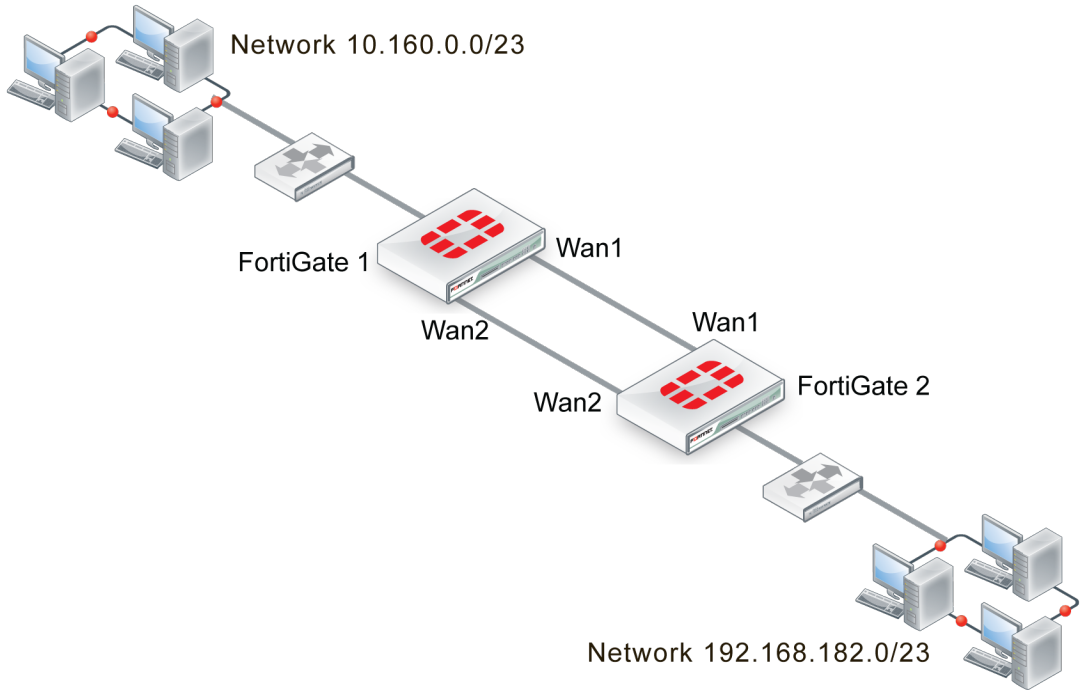

FortiGate 1 should learn the route to network 192.168.182.0 and FortiGate 2 should learn the route to network 10.160.0.0. Under normal conditions, they should learn these routes through the WAN1 link. The WAN2 link should be used only as a backup.

With the default settings, each FortiGate unit learns these routes from both WAN1 and WAN2.

FortiGate 1:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 1 Full/Backup 00:00:33 10.182.0.187 wan1
10.2.2.2 1 Full/Backup 00:00:31 10.183.0.187 wan2
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:01
[110/10] via 10.182.0.187, wan1, 00:00:01
O 192.168.182.0/23 [110/20] via 10.183.0.187, wan2, 00:02:04
[110/20] via 10.182.0.187, wan1, 00:02:04
```

FortiGate 2:

```
FGT2 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.1.1.1 1 Full/DR 00:00:38 10.182.0.57 wan1
10.1.1.1 1 Full/DR 00:00:38 10.183.0.57 wan2
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.183.0.57, wan2, 00:00:39
[110/20] via 10.182.0.57, wan1, 00:00:39
```

Adjusting the route costs

On both FortiGate units, the cost of the route through WAN2 is adjusted higher so that this route will only be used if the route through WAN1 is unavailable. The default cost is 10. The WAN2 route will be changed to a cost of 200.

On both FortiGate units:

```
config router ospf
   config ospf-interface
      edit "WAN2_higher_cost"
         set cost 200
         set interface "wan2"
      end
```

Now, both FortiGate units use only the WAN1 route:

FortiGate 1:

```
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.182.0.187, wan1, 00:00:40
O 192.168.182.0/23 [110/20] via 10.182.0.187, wan1, 00:00:40
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.182.0.57, wan1, 00:09:37
```

LSDB check on FortiGate 1:

```
FGT1 # get router info ospf database router lsa
Router Link States (Area 0.0.0.0)
LS age: 81
Options: 0x2 (*|-|-|-|-|-|E|-)
Flags: 0x0
LS Type: router-LSA
Link State ID: 10.1.1.1
Advertising Router: 10.1.1.1
LS Seq Number: 8000000b
Checksum: 0xe637
Length: 60
Number of Links: 3

   Link connected to: Stub Network
   (Link ID) Network/subnet number: 10.160.0.0
   (Link Data) Network Mask: 255.255.254.0
   Number of TOS metrics: 0
   TOS 0 Metric: 10
   Link connected to: a Transit Network
   (Link ID) Designated Router address: 10.183.0.187
   (Link Data) Router Interface address: 10.183.0.57
   Number of TOS metrics: 0
   TOS 0 Metric: 200

   Link connected to: a Transit Network
   (Link ID) Designated Router address: 10.182.0.57
```

```
      (Link Data) Router Interface address: 10.182.0.57
      Number of TOS metrics: 0
      TOS 0 Metric: 10

   LS age: 83
   Options: 0x2 (*|-|-|-|-|-|E|-)
   Flags: 0x2 : ASBR
   LS Type: router-LSA
   Link State ID: 10.2.2.2
   Advertising Router: 10.2.2.2
   LS Seq Number: 8000000e
   Checksum: 0xfc9b
   Length: 60
      Number of Links: 3

      Link connected to: Stub Network
      (Link ID) Network/subnet number: 192.168.182.0
      (Link Data) Network Mask: 255.255.254.0
      Number of TOS metrics: 0
      TOS 0 Metric: 10

      Link connected to: a Transit Network
      (Link ID) Designated Router address: 10.183.0.187
      (Link Data) Router Interface address: 10.183.0.187
      Number of TOS metrics: 0
      TOS 0 Metric: 200

      Link connected to: a Transit Network
      (Link ID) Designated Router address: 10.182.0.57
      (Link Data) Router Interface address: 10.182.0.187
      Number of TOS metrics: 0
      TOS 0 Metric: 10
```

Verifying route redundancy

Bring down WAN1 and then check the routes on the two FortiGate units.

FortiGate 1:

```
FGT1 # get router info routing-table ospf
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:06
O 192.168.182.0/23 [110/210] via 10.183.0.187, wan2, 00:00:06
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/210] via 10.183.0.57, wan2, 00:00:14
```

The WAN2 interface is now in use on both units.

BGP

This section describes Border Gateway Protocol (BGP).

BGP background and concepts

BGP contains two distinct subsets: internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together and is the main routing protocol for the Internet backbone. FortiGate units support iBGP, and eBGP only for communities.

Background

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in RFC 1771. That RFC has since been replaced by RFC 4271. The main benefits of BGP-4 are classless inter-domain routing and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies, and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in RFC 2858 and RFC 2545.

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP are not explained here because they are common to other dynamic routing protocols. When determining your network topology, note that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. For more information about the parts of BGP that are not listed here, see Dynamic routing overview on page 104.

BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4 but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the "6" on the end of the keyword, such as `config network6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the FortiGate CLI Reference.

IPv6 BGP commands include:

```
config router bgp
   set activate6 {enable | disable}
   set allowas-in6 <max_num_AS_integer>
```

```
    set allowas-in-enable6 {enable | disable}
    set as-override6 {enable | disable}
    set attribute-unchanged6 [as-path] [med] [next-hop]
    set capability-default-originate6 {enable | disable}
    set capability-graceful-restart6 {enable | disable}
    set capability-orf6 {both | none | receive | send}
    set default-originate-route-map6 <routemap_str>
    set distribute-list-in6 <access-list-name_str>
    set distribute-list-out6 <access-list-name_str>
    set filter-list-in6 <aspath-list-name_str>
    set filter-list-out6 <aspath-list-name_str>
    set maximum-prefix6 <prefix_integer>
    set maximum-prefix-threshold6 <percentage_integer>
    set maximum-prefix-warning-only6 {enable | disable}
    set next-hop-self6 {enable | disable}
    set prefix-list-in6 <prefix-list-name_str>
    set prefix-list-out6 <prefix-list-name_str>
    set remove-private-as6 {enable | disable}
    set route-map-in6 <routemap-name_str>
    set route-map-out6 <routemap-name_str>
    set route-reflector-client6 {enable | disable}
    set route-server-client6 {enable | disable}
    set send-community6 {both | disable | extended | standard}
    set soft-reconfiguration6 {enable | disable}
    set unsuppress-map6 <route-map-name_str>
    config network6
    config redistribute6
end
```

Role of routers in BGP networks

Dynamic routing has a number of different roles that routers can fill, such as those covered in Dynamic routing overview on page 104. BGP has a number of custom roles that routers can fill. These include:

- Speaker routers
- Peer routers or neighbors
- Route reflectors

Speaker routers

Any router configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that are not speaker routers are not treated as BGP routers.

Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to your FortiGate unit. Your FortiGate unit learns about all other routers through these peers.

You need to manually configure BGP peers on your FortiGate unit as neighbors. Otherwise, these routers will not be seen as peers, but simply as other routers on the network that do not support BGP. Optionally, you can use MD5 authentication to password-protect BGP sessions with those neighbors (see RFC 2385).

You can configure up to 1000 BGP neighbors on your FortiGate unit. You can clear all or some BGP neighbor connections (sessions), using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:

```
execute router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```

In the following diagram, Router A is directly connected to five other routers in a network that contains 12 routers. These routers (the ones in the blue circle) are Router A's peers or neighbors.

Router A and its five peer routers

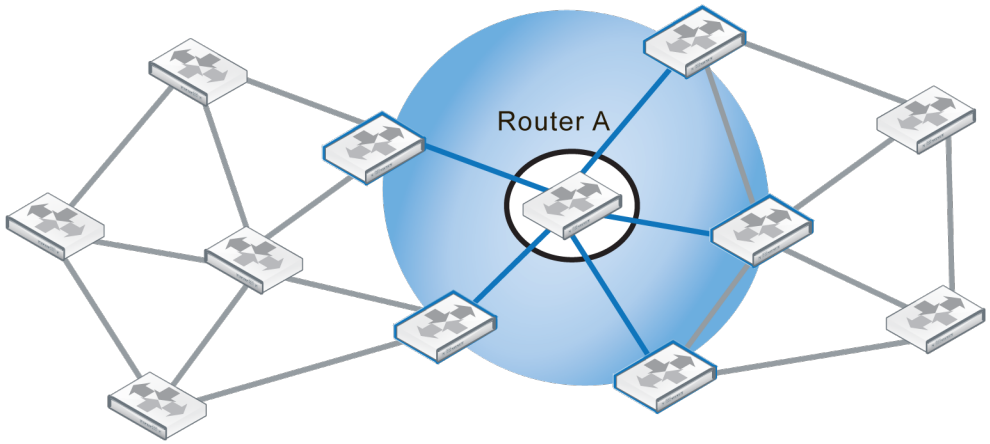

As a minimum, when configuring BGP neighbors, you must enter their IP address and the AS number (remote-as). This is all of the information the web-based manager interface allows you to enter for a neighbor.

The BGP commands related to neighbors are quite extensive and include:

```
config router bgp
   config neighbor
      edit <neighbor_address_ipv4>
         set activate {enable | disable}
         set advertisement-interval <seconds_integer>
         set allowas-in <max_num_AS_integer>
         set allowas-in-enable {enable | disable}
         set as-override {enable | disable}
         set attribute-unchanged [as-path] [med] [next-hop]
         set bfd {enable | disable}
         set capability-default-originate {enable | disable}
         set capability-dynamic {enable | disable}
         set capability-graceful-restart {enable | disable}
         set capability-orf {both | none | receive | send}
         set capability-route-refresh {enable | disable}
         set connect-timer <seconds_integer>
         set description <text_str>
         set distribute-list-in <access-list-name_str>
         set distribute-list-out <access-list-name_str>
         set dont-capability-negotiate {enable | disable}
         set ebgp-enforce-multihop {enable | disable}
```

```
            set ebgp-multihop {enable | disable}
            set ebgp-multihop-ttl <seconds_integer>
            set filter-list-in <aspath-list-name_str>
            set filter-list-out <aspath-list-name_str>
            set holdtime-timer <seconds_integer>
            set interface <interface-name_str>
            set keep-alive-timer <seconds_integer>
            set maximum-prefix <prefix_integer>
            set maximum-prefix-threshold <percentage_integer>
            set maximum-prefix-warning-only {enable | disable}
            set next-hop-self {enable | disable}
            set passive {enable | disable}
            set password <string>
            set prefix-list-in <prefix-list-name_str>
            set prefix-list-out <prefix-list-name_str>
            set remote-as <id_integer>
            set remove-private-as {enable | disable}
            set retain-stale-time <seconds_integer>
            set route-map-in <routemap-name_str>
            set route-map-out <routemap-name_str>
            set route-reflector-client {enable | disable}
            set route-server-client {enable | disable}
            set send-community {both | disable | extended | standard}
            set shutdown {enable | disable}
            set soft-reconfiguration {enable | disable}
            set strict-capability-match {enable | disable}
            set unsuppress-map <route-map-name_str>
            set update-source <interface-name_str>
            set weight <weight_integer>
        end
    end
end
```

Route reflectors

Route reflectors (RR) in BGP concentrate route updates so other routers only need to talk to the RRs to get all of the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. BGP RRs are defined in RFC 1966.

In a BGP RR configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other RRs and border routers. Only the reflectors need to be configured, not the clients, because the clients will find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. FortiGate units can be configured as either reflectors or clients.

Since RRs are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running BGP typically do not require RRs. However, RRs are a useful feature for large companies, where their AS may include 100 routers or more. For example, a full mesh 20 router configuration within an AS, there would have to be 190 unique BGP sessions just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. Based on these numbers, updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how RRs can improve the situation when only six routers are involved. The AS without RRs requires 15 sessions between the routers. In the AS with RRs, the two RRs receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster, as well as other RRs, and pass them on to the border router. The RR configuration requires only six sessions. This example shows a reduction of 60% for the number of required sessions.

Required sessions within an AS with and without RRs

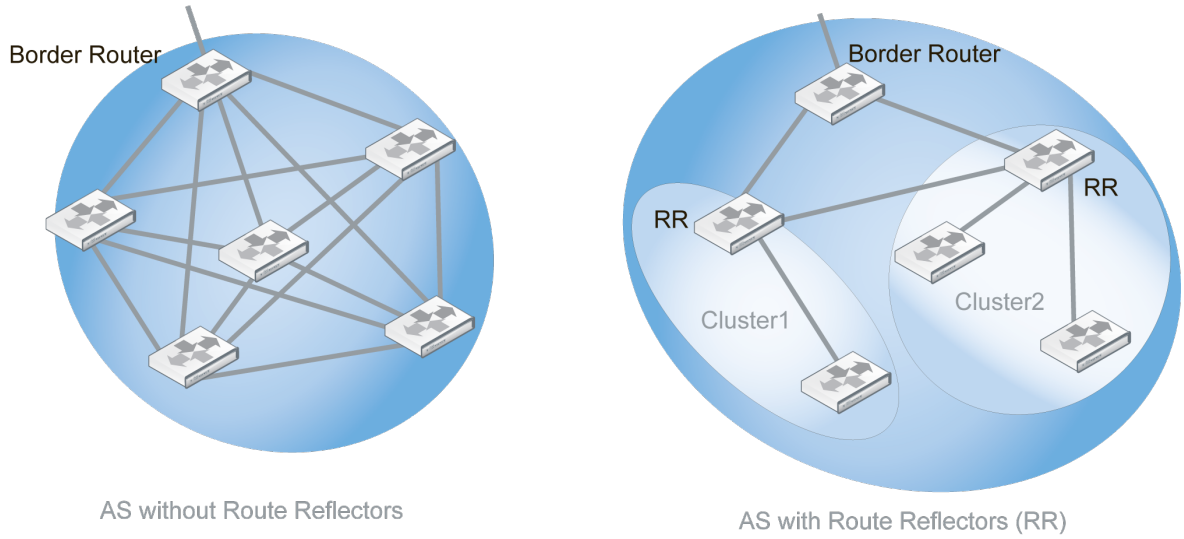

The BGP commands related to RRs include:

```
config router bgp
   config neighbor
      set route-reflector-client {enable | disable}
      set route-server-client {enable | disable}
   end
end
```

Confederations

Confederations were introduced to reduce the number of BGP advertisements on a segment of the network and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units. Confederations are defined in RFC 3065 and RFC 1965.

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications because many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging ASs. Each AS being merged can easily become a confederation, which requires few changes. Any additional permanent changes can then be implemented over time, as required. The diagram below shows the group of ASs before merging and the corresponding confederations afterward, as part of the single AS with the addition of a new border router. It should be noted that after merging, if the border router becomes a route reflector, then each confederation only needs to communicate with one other router instead of five others.

AS merging using confederations

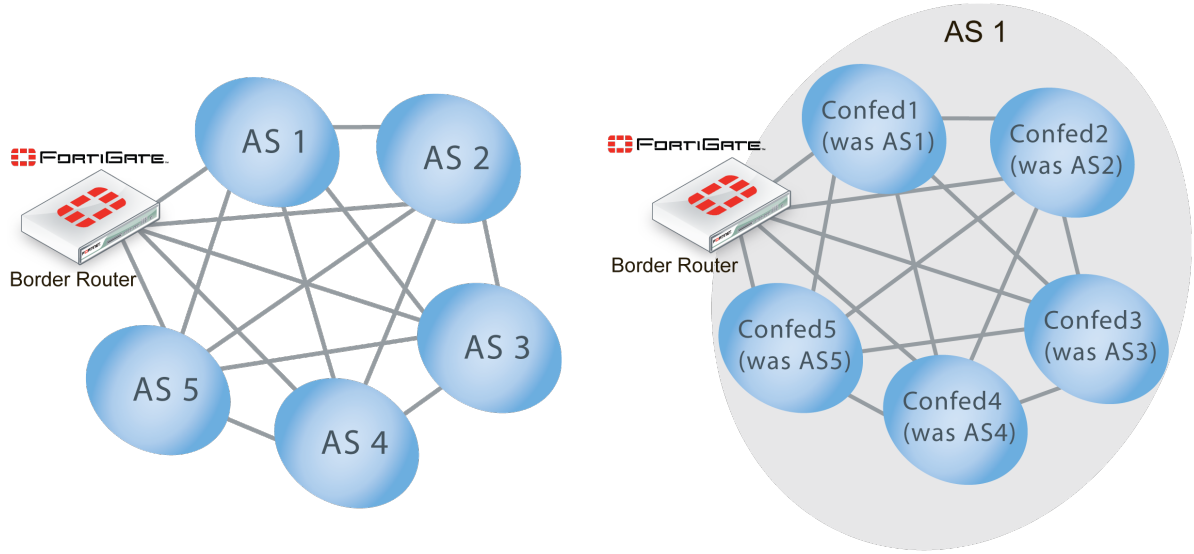

Multiple ASes before merging

Combined AS with confederations and new FortiGate unit border router

Confederations and RRs perform similar functions: they both sub-divide large ASs for more efficient operation. They differ in that route reflector clusters can include routers that are not members of a cluster, whereas routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS_PATH attribute, making it easier to trace.

It is important to note that while confederations essentially create sub-ASs, all the confederations within an AS appear as a single AS to external ASs.

Confederation related BGP commands include:

```
config router bgp
   set confederation-identifier <peerid_integer>
end
```

BGP conditional advertisements

Normally, routes are propagated regardless of the existence of a different path. The BGP conditional advertisement feature allows a route not to be advertised, based on the existence or non-existence of other routes. With this feature, a child table under bgp.neighbor is introduced. Any route matched by one of the route-maps specified in the table will be advertised to the peer, based on the corresponding route-map condition.

You can enable and disable conditional advertisements using the CLI.

To configure BGP conditional advertisements - CLI:

```
config router bgp
   set as 3
      config neighbor
         edit "10.10.10.10"
            set remote-as 3
               config conditional-advertise
                  edit "route-map-to-match-sending"
```

```
                set condition-routemap "route-map-to-match-condition"
                set condition-type [exist | non-exist]
             next
          end
       next
    end
```

BGP neighbor groups

The BGP neighbor group feature allows a large number of neighbors to be configured automatically based on a range of neighbors' source addresses.

To configure BGP neighbor groups - CLI:

Start by adding a BGP neighbor group:

```
config router bgp
   config neighbor-group
      edit <neighbor-group-name>
         set remote-as 100
   ...
```

(All options for BGP neighbor group are supported except `password`.)

```
end
```

Then add a BGP neighbor range:

```
config router bgp
   config neighbor-range
      edit 1
         set prefix 192.168.1.0/24
         set max-neighbor-num 100
         set neighbor-group <neighbor-group-name>
      next
   end
```

Network Layer Reachability Information

Network Layer Reachability Information (NLRI) is unique to BGP-4. It is sent as part of the update messages sent between BGP routers and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that, when combined, are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route and are modified, as required, along the route.

BGP can work well with mostly default settings, but if you are going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include:

Attribute	Description
AS_PATH	A list of ASs a route has passed through. For more information, see AS_PATH on page 206.
MULTI_EXIT_DESC (MED)	Which router to use to exit an AS with more than one external connection. For more information, see MULTI_EXIT_DESC on page 207.
COMMUNITY	Used to apply attributes to a group of routes. For more information, see COMMUNITY on page 207.
NEXT_HOP	Where the IP packets should be forwarded to, like a gateway in static routing. For more information, see NEXT_HOP on page 208.
ATOMIC_ AGGREGATE	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. For more information, see ATOMIC_AGGREGATE on page 208.
ORIGIN	Used to determine if the route is from the local AS or not. For more information, see ORIGIN on page 208.
LOCAL_PREF	Used only within an AS to select the best route to a location (like MED).

Inbound policies on FortiGate units can change the NEXT-HOP,LOCAL-PREF, MED and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the unit cannot affect these attributes.

AS_PATH

AS_PATH is the BGP attribute that keeps track of each AS that a route advertisement has passed through. AS_PATH is used by confederations and by exterior BGP (EBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS_PATH with that router's AS in it. The diagram below shows the route between Router A and Router B. The AS_PATH from A to B would read 701,702,703 for each AS that the route passes through.

As of the beginning of 2010, the industry upgraded from 2-byte to 4-byte AS_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS_PATH numbers. FortiOS supports 4-byte AS_PATHs in its BGP implementation.

AS_PATH of 701,702, 703 between routers A and B

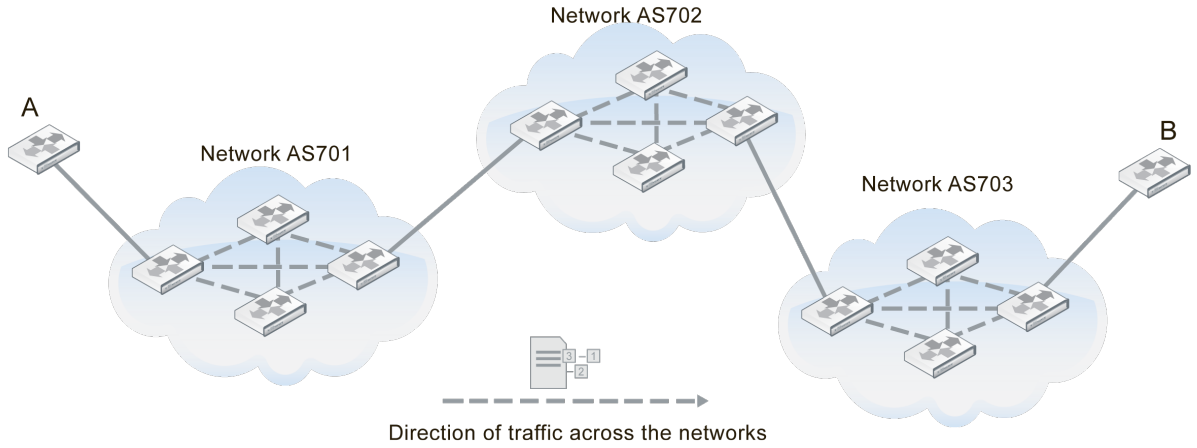

The BGP commands related to AS_PATH include:

```
config router bgp
   set bestpath-as-path-ignore {enable | disable}
end
```

MULTI_EXIT_DESC

BGP AS systems can have one or more routers that connect them to other ASs. For ASs with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes, such as delay. It is a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When the FortiGate unit receives a BGP update, the FortiGate unit examines the MED attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiGate unit routing table.

FortiGate units have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information, which can be suspicious as a possible hacking attempt or an attack on the network. At best, it signifies an unreliable route to select.

The BGP commands related to MED include:

```
config router bgp
   set always-compare-med {enable | disable}
   set bestpath-med-confed {enable | disable}
   set bestpath-med-missing-as-worst {enable | disable}
   set deterministic-med {enable | disable}
   config neighbor
      set attribute-unchanged [as-path] [med] [next-hop]
   end
end
```

COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

The FortiGate unit can set the COMMUNITY attribute of a route to assign the route to predefined paths (see RFC 1997). The FortiGate unit can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include:

```
config router bgp
   set send-community {both | disable | extended | standard}
end
```

NEXT_HOP

The NEXT_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT_HOP attribute is much like a gateway in static routing.

FortiGate units allow you to to change the advertising of the FortiGate unit's IP address (instead of the neighbor's IP address) in the NEXT_HOP information that is sent to IBGP peers. This is changed with the `config neighbor`, `set next-hop-self` command.

The BGP commands related to NEXT_HOP include:

```
config router bgp
   config neighbor
      set attribute-unchanged [as-path] [med] [next-hop]
      set next-hop-self {enable | disable}
   end
end
```

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that is easier to send in updates for. When it reaches its destination, the summarized routes are split back up into the individual routes.

Your FortiGate unit does not specifically set this attribute in the BGP router command, but it is used in the route map command.

The commands related to ATOMIC_AGGREGATE include:

```
config router route-map
   edit <route_map_name>
      config rule
      edit <route_map_rule_id>
         set set-aggregator-as <id_integer>
         set set-aggregator-ip <address_ipv4>
         set set-atomic-aggregate {enable | disable}
      end
   end
end
```

ORIGIN

The ORIGIN attribute records where the route came from. The options can be IBGP, EBGP, or incomplete. This information is important because internal routes (IBGP) are, by default, higher priority than external routes (EBGP). However, incomplete ORIGINs are the lowest priority of the three.

The commands related to ORIGIN include:

```
config router route-map
   edit <route_map_name>
      set comments <string>
      config rule
      edit <route_map_rule_id>
         set match-origin {egp | igp | incomplete | none}
      end
   end
end
```

How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other and establish a connection, they go from the idle state and through the various states until they reach the established state. An error can cause the connection to be dropped and the state of the router to be reset to either active or idle. These errors can be caused by: TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used, such as multiprotocol extensions, that can include IPv6 and VPNs.

IBGP versus EBGP

When you read about BGP, you often see EBGP or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASs) and interior BGP (IBGP) involves packets that stay within a single AS. For example, the AS_PATH attribute is only useful for EBGP where routes pass through multiple ASs.

These two modes are important because some features of BGP are used only for one of EBGP or IBGP. For example, confederations are used in EBGP and RRs are used only in IBGP. Also, routes learned from IBGP have priority over routes learned from EBGP.

FortiGate units have some commands that are specific to EBGP, including:

- automatically resetting the session information to external peers if the connection goes down: `set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (you must also configure local and internal distances if this is set): `set distance-external <distance_integer>`
- enforcing EBGP multihops and their TTL (number of hops): `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

BGP path determination: which route to use

Firstly, recall that the number of available or supported routes is not set by the configuration but depends on your FortiGate's available memory. All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination do not change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute, to allow an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiGate unit receives BGP updates or when the FortiGate unit sends out BGP updates.

The three phases of a BGP routing decision

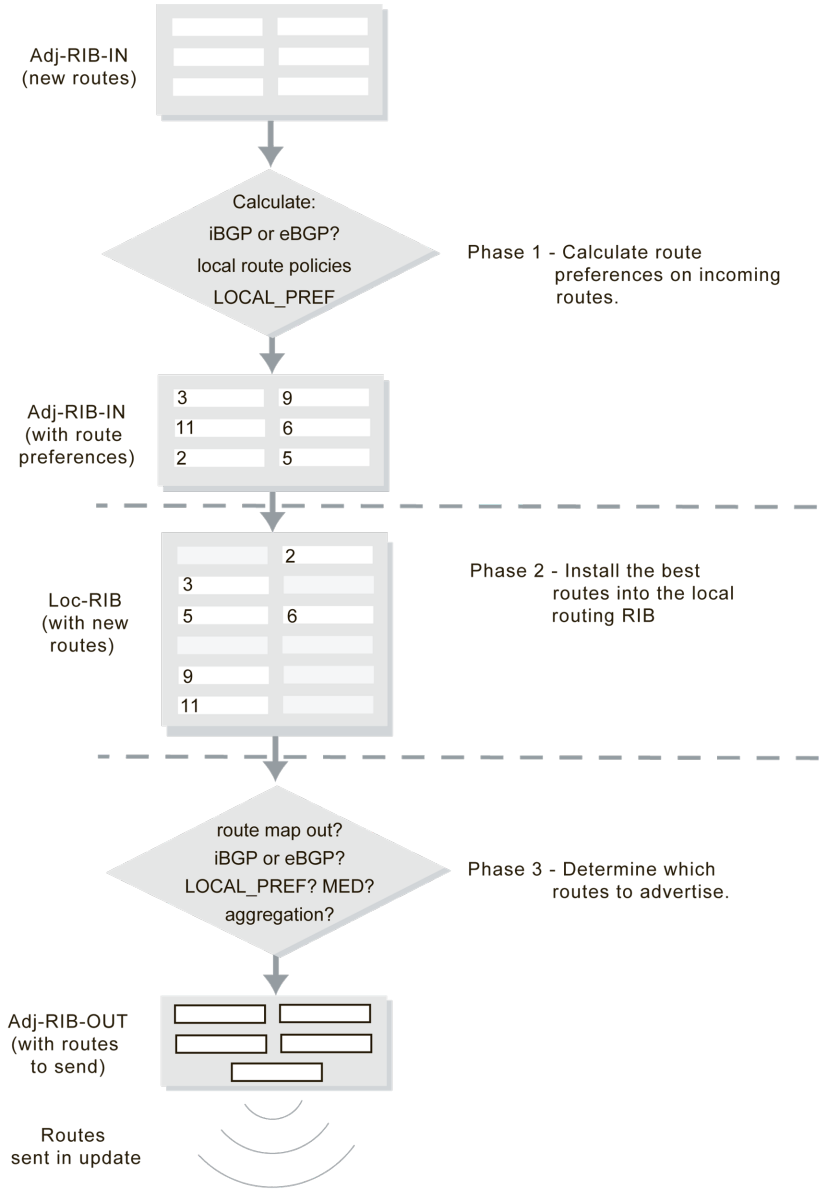

Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRLI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (IBGP), policy information or LOCAL_PREF is used. For external peer learned routes, it is based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the master routing table. Each route from Phase 1 has their NEXT_HOP checked to ensure the destination is reachable. If it is reachable, the AS_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there is only one route to a location, it is installed.
- If there are multiple routes to the same location, use the most preferred route from Level 1.
- If there is a tie, break the tie based on the following, in descending order of importance: shortest AS_PATH, smallest ORIGIN number, smallest MED, EBGP over IBGP, smallest metric or cost for reaching the NEXT_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed, the Loc-RIB will consist of the best of both the new and older routes.

Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there is any route aggregation or summarizing, it happens here. Also, any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

Aggregate routes and addresses

BGP-4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing allows the configuration of aggregate routes by stating the address bits the aggregated addresses have in common. For more information, see Dynamic routing overview on page 104.

The ATOMIC_AGGREGATE attribute informs routers that the route has been aggregated and should not be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are:

```
config router bgp
   config aggregate-address
      edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix <address_ipv4mask>
      set summary-only {enable | disable}
   end
   config aggregate-address6
      edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix6 <address_ipv6mask>
      set summary-only {enable | disable}
   end
```

Configuring BGP graceful restart process on timer

You can configure the BGP graceful restart process to stop only when the restart timer expires, using the following CLI commands:

```
config router bgp
   set graceful-end-on-timer enable
```

Configuring option to bring down BGP neighbor when the link is down

You can configure an option to bring down BGP neighbors when the outgoing interface is down, using the following CLI commands:

```
config router bgp
   config neighbor
      edit <ip_address>
         set linkdown-failover enable
```

Configuring option to keep routes for a period after the BGP neighbor is down

You can configure an option to keep routes for a period after the BGP neighbor is down. If you enable this option for a BGP neighbor, the route learned from the neighbor is kept for the configured `graceful-stalepath-time` after the neighbor is down because of hold timer expiration or TCP connection failure.

To configure this option, use the following CLI commands:

```
config router bgp
   config neighbor
      edit <ip_address>
         set stale-route enable
```

BGP local-AS support

FortiGate supports BGP local-AS. Local-AS allows you to configure a BGP speaker to have a real local-AS and a secondary local-AS for a specific neighbor, so its local-AS number appears different to different neighbors.

You can configure a BGP speaker to have a real local-AS and a secondary local-AS for a specific neighbor, so the local-AS number appears different to neighbor B and neighbor A.

Configuring BGP local-AS

To configure BGP local-AS for BGP peers, use the following CLI commands:

```
config router bgp
   config neighbor
      edit "neighbor" / edit <ip_address>
         …
         set local-as 300 (?) / set local-as <integer>
         set local-as-no-prepend {enable | disable}
         set local-as-replace-as {enable | disable}
      end
```

CLI option	Description
<ip_address>	The IP/IPv6 address of neighbor
<local-as <integer>>	The local-AS number
local-as-no-prepend { enable | disable}	Set this to enable if you do not want to prepent local-AS to incoming updates.
local-as-replace-as { enable | disable}	Set this to enable to replace a real AS with local-AS in outgoing updates.

Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically, the problems with a BGP network that has been configured involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the command:

```
execute router clear bgp as 650001
```

Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term that is used if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables. This creates a lot of administration traffic on the network and the same traffic re-occurs when that router comes back online. If the problem is something like a faulty network cable that wobbles online and offline every 10 seconds, there could easily be an overwhelming amount of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiGate units in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline resulting in route flap. While this does not occur often, or more than once at a time, it can still result in an interruption in traffic which is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they do not expire during the failover process. Also, configuring graceful restart on the HA cluster will help with a smooth failover.

The first method of dealing with route flap should be to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either avoid using that interface or swap in a functioning router. If the power source is bad on a router, either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However, if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include:

- Hold down timer
- Dampening
- Graceful restart
- Bi-directional forwarding detection

Hold down timer

The first line of defense to a flapping route is the hold down timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

Once activated, the hold down timer will not allow the FortiGate unit to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage will be recognized by the FortiGate unit. For the duration of the other outages, there will be no changes because the Fortigate unit is essentially treating this router as down. If the route is still flapping after the timer expires, it will happen all over again.

Even if the route is not flapping (for example, if it goes down, comes up, and stays back up) the timer still counts down and the route is ignored for the duration of the timer. In this situation, the route will be seen as down longer than it really is but there will be only the one set of route updates. This is not a problem in normal operation because updates are not frequent.

Also, the potential for a route to be treated as down when it is really up can be viewed as a robustness feature. Typically, you do not want most of your traffic being routed over an unreliable route. So if there is route flap going on, it is best to avoid that route if you can. This is enforced by the hold down timer.

How to configure the hold down timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the hold down timer.

For example, your network has two routes that you want to set the hold down timer for. One is your main route (to 10.12.101.4) that all of your Internet traffic goes through, and it cannot be down for long if it is down. The second is a low speed connection to a custom network that is used infrequently (to 10.13.101.4). The hold down timer for the main route should be fairly short, let us say 60 seconds instead of the default 180 seconds. The second route timer can be left at the default, or even longer since it is rarely used. In your BGP configuration this looks like:

```
config router bgp
   config neighbor
   edit 10.12.101.4
      set holddown-timer 60
   next
   edit 10.13.101.4
      set holddown-timer 180
   next
   end
end
```

Dampening

Dampening is a method used to limit the amount of network problems due to flapping routes. With dampening, the flapping still occurs but the peer routers pay less and less attention to that route as it flaps more often. One flap does not start dampening, but the second flap starts a timer where the router will not use that route because it is considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There is a period of time called the reachability half-life, after which a route flap will be suppressed for only half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiGate units cache by using one of the `execute router clear bgp` commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
```
or
```
execute router clear bgp flap-statistics {<ip> | <ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the command:

```
execute router clear bgp dampening 10.10.0.0/16
```

The BGP commands related to route dampening are:

```
config router bgp
   set dampening {enable | disable}
   set dampening-max-suppress-time <minutes_integer>
   set dampening-reachability-half-life <minutes_integer>
   set dampening-reuse <reuse_integer>
   set dampening-route-map <routemap-name_str>
   set dampening-suppress <limit_integer>
   set dampening-unreachability-half-life <minutes_integer>
end
```

Graceful restart

BGP4 has the capability to gracefully restart.

In some situations, route flap is caused by routers that appear to be offline but the hardware portion of the router (control plane) can continue to function normally. One example of this is when some software is restarting or being upgraded but the hardware can still function normally.

Graceful restart is best used for these situations where routing will not be interrupted, but the router is unresponsive to routing update advertisements. Graceful restart does not have to be supported by all routers in a network, but the network will benefit when more routers support it.

FortiGate HA clusters can benefit from graceful restart. When a failover takes place, the HA cluster will advertise that it is going offline, and will not appear as a route flap. It will also enable the new HA main unit to come online with an updated and usable routing table. If there is a flap, the HA cluster routing table will be out-of-date.

For example, your FortiGate unit is one of four BGP routers that send updates to each other. Any of those routers may support graceful starting. When a router plans to go offline, it sends a message to its neighbours stating how long it expects to be offline. This way, its neighboring routers do not remove it from their routing tables. However, if that router is not back online when expected, the routers will mark it offline. This prevents routing flap and its associated problems.

Scheduled time offline

Graceful restart is a means for a router to advertise that it is going to have a scheduled shutdown for a very short period of time. When neighboring routers receive this notice, they will not remove that router from their routing table until after a set time elapses. During that time, if the router comes back online, everything continues to function as normal. If that router remains offline longer than expected, then the neighboring routers will update their routing tables as they assume that router will be offline for a long time.

FortiGate units support both graceful restart of their own BGP routing software and neighboring BGP routers.

For example, if a neighbor of your FortiGate unit with an IP address of 172.20.120.120 supports graceful restart, enter the command:

```
config router bgp
   config neighbor
   edit 172.20.120.120
      set capability-graceful-restart enable
   end
end
```

If you want to configure graceful restart on your FortiGate unit where you expect the Fortigate unit to be offline for no more than 2 minutes, and after 3 minutes the BGP network should consider the FortiGate unit to be offline, enter the command:

```
config router bgp
   set graceful-restart enable
   set graceful-restart-time 120
   set graceful-stalepath-time 180
end
```

The BGP commands related to BGP graceful restart are:

```
config router bgp
   set graceful-restart { disable| enable}
   set graceful-restart-time <seconds_integer>
   set graceful-stalepath-time <seconds_integer>
   set graceful-update-delay <seconds_integer>
   config neighbor
      set capability-graceful-restart {enable | disable}
   end
end

execute router restart
```

Before the restart, the router sends its peers a message to say it is restarting. The peers mark all the restarting router's routes as stale, but they continue to use the routes. The peers assume the router will restart, check its routes, and take care of them, if needed, after the restart is complete. The peers also know what services the restarting router can maintain during its restart. After the router completes the restart, the router sends its peers a message to say it is done restarting.

Bi-directional forwarding detection

Bi-directional Forwarding Detection (BFD) is a protocol used to quickly locate hardware failures in the network. Routers running BFD communicate with each other, and if a timer runs out on a connection, that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated.

While BGP can detect route failures, BFD can be configured to detect these failures more quickly which allows for faster responses and improved convergence. This can be balanced with the bandwidth BFD uses in its frequent route checking.

Configurable granularity

BFD can run on the entire FortiGate unit, selected interfaces, or on BGP for all configured interfaces. The hierarchy allows each lower level to override the upper level's BFD setting. For example, if BFD was enabled for the FortiGate unit, it could be disabled only for a single interface or for BGP. For information about FortiGate-wide BFD options, see config system settings in the FortiGate CLI Reference.

BFD can be configured only through the CLI.

The BGP commands related to BFD are:

```
config system {setting | interface}
   set bfd {enable | disable | global}
   set bfd-desired-mix-tx <milliseconds>
   set bfd-detect-mult <multiplier>
   set bfd-required-mix-rx <milliseconds>
   set bfd-dont-enforce-src-port {enable | disable}

config router bgp
   config neighbor
      edit <neighbor_address_ipv4>
      set bfd {enable | disable}
   end
end

get router info bfd neighbor
execute router clear bfd session <src_ipv4> <dst_ipv4> <interface>
```

The `config system` commands allow you to configure whether BFD is enabled in a particular unit/vdom or individual interface, and how often the interface requires the sending and receiving of BFD information.

The `config router bgp` commands allow you to set the addresses of the neighbor units that are also running BFD. Both units must be configured with BFD in order to make use of it.

Dual-homed BGP example

This is an example of a small network that uses BGP routing connections to two ISPs. This is a common configuration for companies that need redundant connections to the Internet for their business.

This configuration is for a small company connected to two ISPs. The company has one main office, the Head Office, and uses static routing for internal routing on that network.

Both ISPs use BGP routing and connect to the Internet directly. They want the company to connect to the ISP networks using BGP. They also use graceful restart to prevent updates that are not needed and use smaller timer values to detect network failures faster.

As can be expected, the company wants to keep their BGP configuration relatively simple and easy to manage. The current configuration has only 3 routers to worry about: the 2 ISP border routers and the FortiGate unit. This means that the FortiGate unit will have only two neighbor routers to configure.

This configuration has the added benefit of being easy to expand if the company wants to add a remote office in the future.

To keep the configuration simple, the company is allowing only HTTP, HTTPS, FTP, and DNS traffic out of the local network. This will allow employees access to the Internet and their web mail.

Why dual home?

Dual homing means having two separate independent connections to the Internet. Servers in this configuration have also been called bastion hosts and can include DNS servers which require multiple connections.

Benefits of dual homing can include:

- Redundant Internet connection that essentially never fails
- Faster connections through one ISP or the other for some destinations, such as other clients of those ISPs

- Load balancing traffic to the company network
- Easier to enable more traffic through two connections than upgrading one connection to bigger bandwidth
- Easier to create protection policies for different traffic through a specific ISP

Some companies require reliable Internet access at all times as part of their business. Consider a doctor operating remotely who has their Internet connection fail — the consequences can easily be life or death.

Dual homing is an extra expense for the second ISP connection and more work to configure and maintain the more complex network topology.

Potential dual homing issues

BGP comes with load balancing issues and dual homing is in the same category. BGP does not inherently deal well with load balancing or getting default routes through BGP. Ideally, one connection may be best for certain destinations but it may not have that traffic routed to it, which makes the load balancing less than perfect. This kind of fine tuning can be very time consuming and usually results in a best effort situation.

When dual homing is not configured properly, your network may become a link between your ISPs and result in very high traffic between the ISPs that does not originate from your network. The problem with this situation is that your traffic may not have the bandwidth it needs, and you will also be paying for a large volume of traffic that is not yours. This problem can be solved by not broadcasting or redistributing BGP routes between the ISPs.

If you learn your default routes from the ISPs, in this example, you may run into an asymmetric routing problem where your traffic loops out one ISP and back to you through the other ISP. If you think this may be happening, you can turn on asymmetric routing on the FortiGate unit (`config system settings, set asymmetric enable`) to verify if that is the problem. Turn this feature off once this is established, since it disables many features on the FortiGate by disabling stateful inspection. Solutions to this problem can include using static routes for default routes instead of learning them through BGP or configuring VDOMs on your FortiGate unit to provide a slightly different path back that is not a true loop.

Network layout and assumptions

The network layout for the basic BGP example involves the company network being connected to both ISPs as shown below. In this configuration, the FortiGate unit is the BGP border router between the Company AS, ISP1's AS, and ISP2's AS.

The components of the layout include:

- The Company AS (AS number 1) is connected to ISP1 and ISP2 through the FortiGate unit.
- The Company has one internal network: the Head Office network at 10.11.101.0/24.
- The FortiGate unit internal interface is on the the company internal network with an IP address of 10.11.101.110.
- The FortiGate unit external1 interface is connected to ISP1's network with an IP address of 172.20.111.5, which is an address supplied by the ISP.
- The FortiGate unit external2 interface is connected to IPS2's network with an IP address of 172.20.222.5, which is an address supplied by the ISP.
- ISP1 AS has an AS number of 650001 and ISP2 has an AS number of 650002.
- Both ISPs are connected to the Internet.
- The ISP1 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.21.111.4.
- The ISP2 border router is a neighbor (peer) of the FortiGate unit. It has an address of 172.22.222.4.
- Apart from graceful restart and shorter timers (holdtimer and keepalive), default settings are to be used whenever possible.

Basic BGP network topology

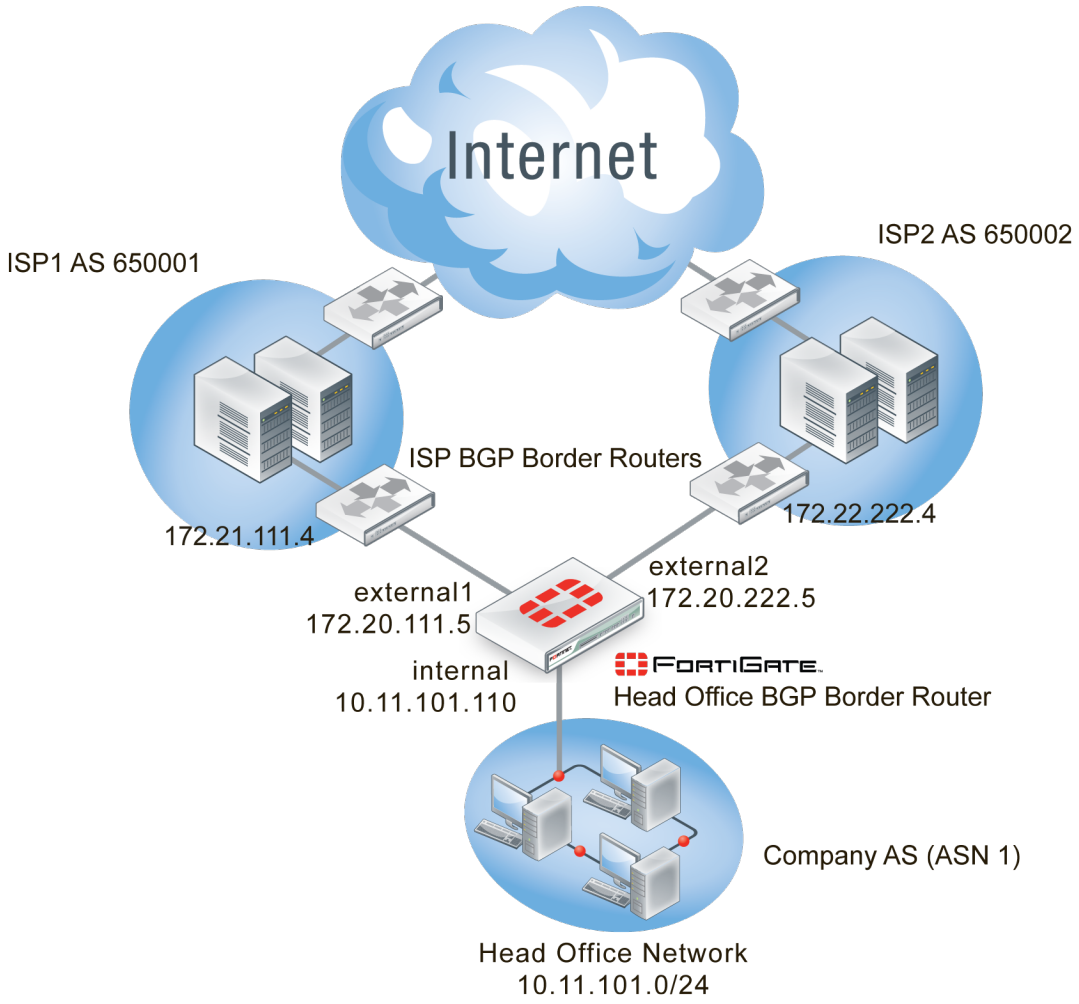

Assumptions

The basic BGP configuration procedure follows these assumptions:

- ISP1 is the preferred route and ISP2 is the secondary route
- All basic configuration can be completed in both the GUI and CLI
- Only one AS is used for the company

For these reasons, this example configuration does not include:

- Bi-directional forwarding detection (BFD)
- Route maps
- Access lists
- Changing redistribution defaults (make link when example is set up)
- IPv6

For more information about these features, see the corresponding section.

Configuring the FortiGate unit

In this topology, the FortiGate unit is the link between the company network and the ISP network. The FortiGate unit is the only BGP router on the company network, but there is at least one other BGP router on the ISP network. There may be more BGP routers, but we do not have that information.

As mentioned in the general configuration steps, the ISP must be notified of the company's BGP router configuration when complete as it will need to add the FortiGate BGP router as a neighbor router on its domain. This step is required for the FortiGate unit to receive BGP routing updates from the ISP network and outside networks.

If the ISP has any special BGP features enabled, such as graceful restart or route dampening, that should be determined ahead of time so those features can be enabled on the FortiGate unit.

To configure the FortiGate unit as a BGP router

1. Configure interfaces and default routes
2. Configure firewall services, addresses, and policies
3. Set the FortiGate BGP information
4. Add the internal network to the AS
5. Additional FortiGate BGP configuration

Configure interfaces and default routes

The FortiGate unit is connected to three networks: the company network on the internal interface, the ISP1 network on the external1 interface, and the ISP2 network on the external2 interface.

This example uses basic interface settings. Check with your ISP to determine if additional settings are required, such as setting the maximum MTU size or if gateway detection is supported.

High end FortiGate units do not have interfaces labeled as Internal or External. Instead, for clarity, we are using the alias feature to name interfaces for these roles.

Default routes to both external interfaces are configured here also. Both are needed in case one goes offline. ISP1 is the primary connection and has a smaller administrative distance so it will be preferred over ISP2. Both distances are set low so they will be preferred over any learned routes.

To configure the FortiGate interfaces - web-based manager

1. Go to **Network > Interfaces**.
2. Edit port 1 (internal) interface.
3. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Company internal network
Interface State	Enabled

4. Edit port 2 (external1) interface.
5. Set the following information and select **OK**.

Alias	external1
IP/Network Mask	172.21.111.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	ISP1 External BGP network
Interface State	Enabled

6. Edit port 3 (external2) interface.
7. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	172.22.222.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	ISP2 External BGP network
Interface State	Enabled

To configure the FortiGate interfaces - CLI

```
config system interface
   edit port1
      set alias internal
      set ip 10.11.101.110 255.255.255.0
      set allowaccess http https ssh
      set description “Company internal network”
      set status up
   next
   edit port2
      set alias external1
      set ip 172.21.111.5 255.255.255.0
      set allowaccess https ssh
      set description “ISP1 External BGP network”
      set status up
   next
   edit port3
      set alias external2
      set ip 172.22.222.5 255.255.255.0
      set allowaccess https ssh
      set description “ISP2 External BGP network”
      set status up
   next
end
```

To configure default routes for both ISPs - web-based manager

1. Go to **Network > Static Routes**.
2. Delete any existing routes with a IP/Mask of address of 0.0.0.0/0.0.0.0
3. Select **Create New** and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.21.111.5
Interface	port2
Administrative Distance	10

4. Select **OK**.
5. Select **Create New** and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.22.222.5
Interface	port3
Administrative Distance	15

6. Select **OK**.

To configure default routes for both ISPs - CLI

```
config router static
   edit 1
      set device "port2"
      set distance 10
      set gateway 172.21.111.5
   next
   edit 2
      set device "port3"
      set distance 15
      set gateway 172.22.222.5
   next
end
```

Configure firewall services, addresses, and policies

To create the security policies, you create the firewall services group that will include all the services that will be allowed, define the addresses that will be used in the security policies, and configure the security policies themselves.

To keep the configuration simple, the company is allowing only HTTP traffic out of the local network. This will allow employees access to the Internet and their web mail. DNS services will also be allowed through the firewall.

The security policies will allow HTTP traffic (port 80 and port 8080), HTTPS traffic (port 443), FTP traffic (port 21), and DNS traffic (port 53 and port 953) in both directions. Also, BGP (port 179) may need access through the firewall.

For added security, you may want to define a smaller range of addresses for the internal network. For example, if only 20 addresses are used, only allow those addresses in the range.

To keep things simple, a zone will be used to group the two ISP interfaces together. This will allow for the use of one security policy to apply to both ISPs at the same time. Remember to block intra-zone traffic as this will help prevent one ISP sending traffic to the other ISP through your FortiGate unit using your bandwidth. The zone keeps configuration simple and if there is a need for separate policies for each ISP in the future, they can be created and the zone can be deleted.

The addresses that will be used are the addresses of the FortiGate unit internal and external ports and the internal network.

More policies or services can be added in the future as applications are added to the network. For more information about security policies, see the firewall chapter in the FortiGate Administration Guide.

When configuring security policies, always enable logging to help you track and debug your traffic flow.

To create a firewall services group - web-based manager

1. Go to **Policy & Objects > Services**, select the dropdown arrow next to **Create New** and select **Service Group**.
2. For **Group Name**, enter “Basic_Services”.
3. From the **Members** dropdown, choose the following six services: BGP, FTP, FTP_GET, FTP_PUT, DNS, HTTP, and HTTPS.
4. Select **OK**.

To create a firewall services group - CLI

```
config firewall service group
   edit "Basic_Services"
      set member "BGP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
   next
end
```

To create a zone for the ISP interfaces - web-based manager

1. Go to **Network > Interfaces**.
2. Select the caret to the right of **Create New** and then select **Zone**.
3. Enter the following information:

Name	ISPs
Block intra-zone traffic	enable
Interface Members	port2 port3

4. Select **OK**.

To create a zone for the ISP interfaces - CLI

```
config system zone
   edit "ISPs"
      set interface "port2" "port3"
      set intrazone block
   next
end
```

To add the firewall addresses - web-based manager

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** and set the following information.

Category	Address
Name	Internal_network
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	port1

3. Select **OK**.

To add the firewall addresses - CLI

```
config firewall address
   edit "Internal_network"
      set associated-interface "port1"
      set subnet 10.11.101.0 255.255.255.0
   next
end
```

To add the HTTP and DNS security policies - web-based manager

1. Go to **Policy & Objects > IPv4 Policy**, and select **Create New**.
2. Set the following information.

Incoming Interface	port1(internal)
Outgoing Interface	ISPs
Source	Internal_network
Destination	All
Schedule	Always

Service	Basic_services
Action	ACCEPT
Firewall / Network Options	Enable NAT
Log Allowed Traffic	Enable
Comments	ISP1 basic services out policy

3. Select **OK**.
4. Select **Create New** and set the following information.

Incoming Interface	ISPs
Outgoing Interface	port1(internal)
Source	All
Destination	Internal_network
Schedule	Always
Service	Basic_services
Action	ACCEPT
Firewall / Network Options	Enable NAT
Log Allowed Traffic	Enable
Comments	ISP1 basic services in policy

To add the security policies - CLI

```
config firewall policy
   edit 1
      set srcintf "port1"
      set srcaddr "Internal_network"
      set dstintf "ISPs"
      set dstaddr "all"
      set schedule "always"
      set service "Basic_services"
      set action accept
      set nat enable
      set profile-status enable
      set logtraffic enable
      set comments "ISP1 basic services out policy"
   next
   edit 2
      set srcintf "ISPs"
      set srcaddr "all"
```

```
        set dstintf "port1"
        set dstaddr "Internal_network"
        set schedule "always"
        set service "Basic_services"
        set action accept
        set nat enable
        set profile-status enable
        set logtraffic enable
        set comments "ISP1 basic services in policy"
    next
end
```

Set the FortiGate BGP information

When using the default information, there are only two fields to set to configure the FortiGate unit as a BGP router.

For this configuration, the FortiGate unit will be in a stub area with one route out — the ISP BGP router. Until you configure the ISP router as a neighbor, even that route out is not available. So, while after this part of the configuration is complete, your FortiGate unit will be running BGP, it will not know about any other routers running BGP until the next part of the configuration is complete.

To set the BGP router information - web-based mananger

1. Go to **Network > BGP**.
2. Set the following information and select **OK**.

Local AS	1
Router ID	10.11.101.110

To set the BGP router information - CLI

```
config router BGP
   set as 1
   set router-id 10.11.101.110
end
```

Add the internal network to the AS

The company is one AS with the FortiGate unit configured as the BGP border router connecting that AS to the two ISPs ASs. The internal network in the Company's AS must be defined. If there were other networks in the company, such as regional offices, they would be added here as well.

To set the networks in the AS - web-based manager

1. Go to **Network > BGP**.
2. Under **Networks**, set the **IP/Netmask** to `10.11.101.0/255.255.255.0` and select **Add**.

To set the networks in the AS - CLI

```
config router bgp
   config network
   edit 1
      set prefix 10.11.101.0 255.255.255.0
```

```
    next
    end
end
```

Add BGP neighbor information

The configuration will not work unless you set **Remote AS** neighbors. This can be done in either the web-based manager or the CLI.

To configure the BGP neighbors - web-based manager

1. Go to **Network > BGP**.
2. Add a **Neighbors IP** of 172.21.111.4 with the **Remote AS** set to 650001, then click **Add/Edit**.
3. Add another **Neighbors IP** of 172.22.222.4 with the **Remote AS** set to 650002, then click **Add/Edit**.

To configure the BGP neighbors - CLI

```
config router BGP
   set as 1
      config neighbor
         edit "172.21.111.4"
            set remote-as 650001
         next
         edit "172.22.222.4"
            set remote-as 650002
         next
      end
   end
```

Additional FortiGate BGP configuration

At this point, those are all the settings that can be done in both the web-based manger and the CLI. The remaining configuration must be completed in the CLI.

These additional settings are mainly determined by your ISP requirements. They will determine your timers, such as keepalive timers, if extended features like BFD and graceful restart are being used, and so on. For this example, some common simple features are being used to promote faster detections of network failures, which will result in better service for the company's internal network users.

The ISPs do not require authentication between peer routers.

These commands will enable or modify the following features on the FortiGate unit and, where possible, on neighboring routers as well:

- `bestpath-med-missing-as-worst`: treats a route without an MED as the worst possible available route due to expected unreliability
- `fast-external-failover`: immediately reset the session information associated with BGP external peers if the link used to reach them goes down
- `graceful-restart*`: advertise reboots to neighbors so they do not see the router as offline, wait before declaring them offline, and how long to wait when they reboot before advertising updates. These commands apply to neighbors and are part of the BGP capabilities. This prevents unneeded routing updates.
- `holdtime-timer`: how long the router will wait for a keepalive message before declaring a router offline. A shorter time will find an offline router faster.

- `keepalive-timer`: how often the router sends out keepalive messages to neighbor routers to maintain those sessions.
- `log-neighbor-changes`: log changes to the status of neighbor routers. This can be useful for troubleshooting from both internal and external networks.
- `connect-timer`: how long (in seconds) the FortiGate unit will try to reach this neighbor before declaring it offline.
- `weight`: used to prefer routes from one neighbor over the other. In this example, ISP1 is the primary connection so it is weighted higher than ISP2

To configure additional BGP options - CLI

```
config router bgp
   set bestpath-med-missing-as-worst enable
   set fast-external-failover enable
   set graceful-restart enable
   set graceful-restart-time 120
   set graceful-stalepath-time 180
   set graceful-update-delay 180
   set holdtime-timer 120
   set keepalive-timer 45
   set log-neighbor-changes enable
   config neighbor
      edit 172.21.111.4
         set connect-timer 60
         set description "ISP1"
         set holdtime-timer 120
         set keepalive-timer 45
         set weight 250
      next
      edit 172.22.222.4
         set connect-timer 60
         set description "ISP2"
         set holdtime-timer 120
         set keepalive-timer 45
         set weight 100
      next
   end
end
```

Configuring other networking devices

There are two other networking devices that need to be configured: the BGP routers for both ISPs.

The ISPs' routers must add the FortiGate unit as a neighbor so route updates can be sent in both directions. Note that ISP1 is not directly connected to ISP2, that we are aware of.

Inform both of your ISPs of your FortiGate unit's BGP information. Once they have configured their router, you can test your BGP connection to the Internet.

They will require your FortiGate unit's:

- IP address of the connected interface
- The router ID
- Your company's AS number

Testing this configuration

With the dual-homed BGP configuration in place, you should be able to send and receive traffic, send and receive routes, and not have any routing loops. Testing the networks will confirm that things are working as expected.

In general, for routing, you need to look at the routing table on different routers to see what routes are being installed. You also need to sniff packets to see how traffic is being routed in real-time. These two sources of information will normally tell you what you need to know.

Testing of this example's network configuration should be completed in the following parts:

- Testing network connectivity
- Verifying the FortiGate unit's routing tables
- Verifying traffic routing
- Verifying the dual-homed side of the configuration

Testing network connectivity

A common first step in testing a new network topology is to test to see if you can reach the Internet and other locations as expected. If not, you may be prevented by cabling issues, software, or other issues.

The easiest way to test connections is to use ping, once you ensure that all the FortiGate unit's interfaces and ISP routers have ping support enabled. Also, ensure that the security policies allow ping through the firewall.

Connections to test, in this example, are the internal network to ISP1's router or the Internet, and the same for ISP2. If you can connect on the external side of the Fortinet unit, try to ping the internal network. These three tests should prove your basic network connections are working.

Once you have finished testing the network connectivity, turn off ping support on the external interfaces for additional security.

Verifying the FortiGate unit's routing tables

The FortiGate routing table contains the routes that are stored for future use. If you are expecting certain routes to be there and they are not, this is a good indicator that your configuration is not what you expected.

The `get router info routing-table details` CLI command will provide you with the routing protocol, destination address, gateway address, interface, and weighting for every route, as well as if the address is directly connected or not.

If you want to limit the display to BGP routes only, use the `get router info routing-table bgp` CLI command. If there are no BGP routes in the routing table, nothing will be displayed. In the CLI command, you can replace BGP with static, or other routing protocols, to only display those routes.

If you want to see the contents of the routing information database (RIB), use the `get router info routing-table database` CLI command. This will display the incoming routes that may or may not make it into the routing table.

Verifying traffic routing

Traffic may be reaching the internal network, but it may be using a different route than you think to get there.

Use a browser to try to access the Internet.

If needed, allow traceroute and other diag ports to be opened until things are working properly. Then remove access for them again.

Look for slow hops on the traceroute, or pings to a location, as they may indicate network loops that need to be fixed.

Any locations that have an unresolved traceroute or ping must be examined and fixed.

Use network packet sniffing to ensure traffic is being routed as you expect.

Verifying the dual-homed side of the configuration

Since there are two connections to the Internet in this example, theoretically you can pull the plug on one of the ISP connections, and all traffic will go through the other connection. Alternately, you may choose to remove a default route to one ISP, remove that ISP's neighbor settings, or change the weightings to prefer the other ISP. These alternate ways to test dual-homing do not change physical cabling, which may be preferred in some situations.

If this does not work as expected, things to check include:

- Default static routes: If these are wrong or do not exist, the traffic cannot get out.
- BGP neighbor information: If the ISP router information is incorrect, the FortiGate unit will not be able to talk to it.

Redistributing and blocking routes in BGP

During normal BGP operation, peer routers redistribute routes from each other. However, in some specific situations it may be best not to advertise routes from one peer, such as if the peer is redundant with another peer (they share the same routes exactly), if it might be unreliable in some way or for some other reason. The FortiGate can also take routes it learns from other protocols and advertise them in BGP, for example OSPF or RIP. If your company hosts its own web or email servers, external locations will require routes to your networks to reach those services.

In this example, the company has an internal network in an OSPF area and is connected to a BGP AS and two BGP peers. The company goes through these two peers to reach the Internet. However, Peer 1 routes will not be advertised to Peer 2. The company internal user and server networks are running OSPF, and will redistribute those routes to BGP so external locations can reach the web and email servers.

Network layout and assumptions

The network layout for the BGP redistributing routes example involves the company network being connected to two BGP peers, as shown below. In this configuration, the FortiGate unit is the BGP border router between the Company AS and the peer routers.

The components of the layout include:

- There is only one BGP AS in this example shared by the FortiGate unit and both peers: AS 65001.
- The company's FortiGate unit connects to the Internet through two BGP peers.
- The company's internal networks on the dmz interface of the FortiGate unit with an IP of 10.11.201.0/24.
- The FortiGate unit's interfaces are connected as follows:
  - port1 (dmz) has IP 10.11.201.110 and is the internal user and server network
  - port2 (external1) has IP 172.21.111.4 and is connected to Peer 1's network

  - port3 (external2) has IP 172.22.222.4 and is connected to Peer 2's network
- Peer 1 has IP 172.21.111.5, and Peer 2 has IP 172.22.222.5.
- OSPF Area 1 is configured on the dmz interface of the FortiGate unit, and is the routing protocol used by the internal users and servers.

BGP network topology

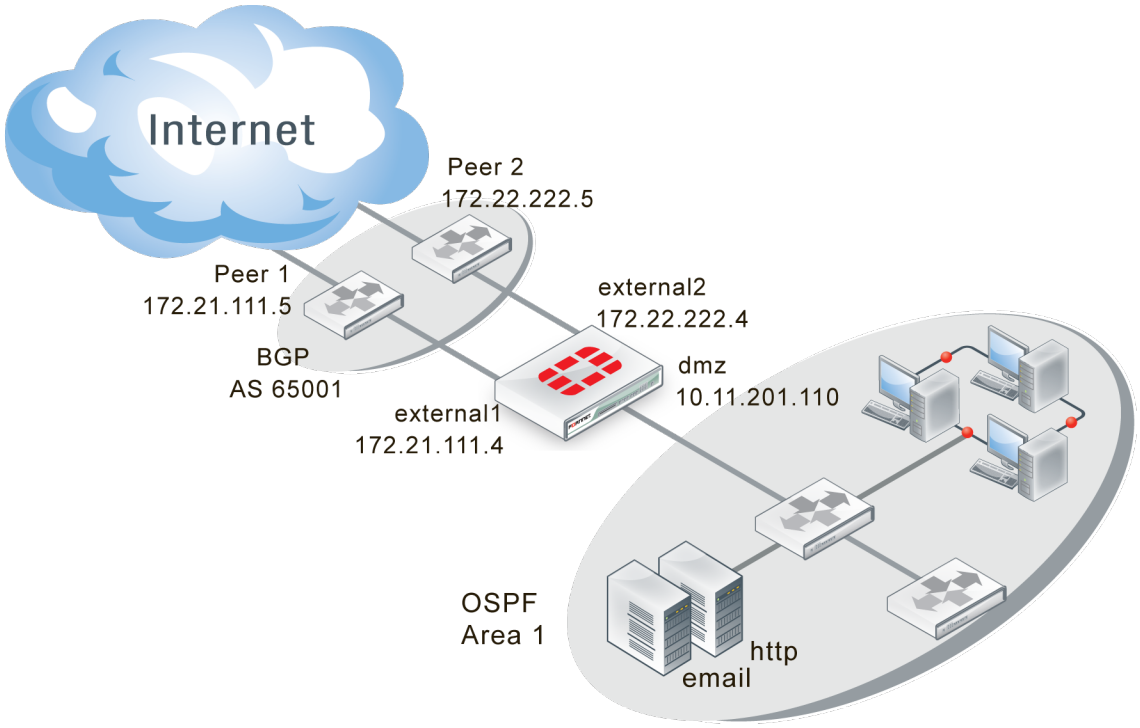

Assumptions

The BGP redistributing routes configuration procedure follows these assumptions:

- The FortiGate unit has been configured following the Install Guide
- Interfaces port1, port2, and port3 exist on the FortiGate unit
- We do not know the router manufacturers of Peer 1 and Peer 2
- We do not know what other devices are on the BGP AS or OSPF Area
- All basic configuration can be completed in both GUI and CLI
- Access lists and route maps will only be configured in CLI
- VDOMs are not enabled on the FortiGate unit

Configuring the FortiGate unit

1. Configuring networks and firewalls on the FortiGate unit
2. Configuring BGP on the FortiGate unit
3. Configuring OSPF on the FortiGate unit
4. Configuring other networking devices
5. Configuring ECMP support for BGP

Configuring networks and firewalls on the FortiGate unit

The FortiGate unit has three interfaces connected to networks: two external and one dmz.

Security policies must be in place to allow traffic to flow between these networks.

Firewall services will change depending on which routing protocol is being used on that network: either BGP or OSPF. Beyond that, all services that are allowed will be allowed in both directions due to the internal servers. The services allowed are web server services (DNS, HTTP, HTTPS, SSH, NTP, FTP*, SYSLOG, and MYSQL), email services (POP3, IMAP, and SMTP), and general troubleshooting services (PING, TRACEROUTE). To increase security, PING and TRACEROUTE can be removed once the network is up and working properly. Other services can be added later, as needed.

To configure the interfaces - web-based manager

1. Go to **Network > Interfaces**.
2. Edit port1 (dmz) interface.
3. Set the following information and select **OK**.

Alias	dmz
IP/Network Mask	10.11.201.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	OSPF internal networks
Administrative Status	Up

4. Edit port2 (external1) interface.
5. Set the following information and select **OK**.

Alias	external1
IP/Network Mask	172.21.111.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external Peer 1
Administrative Status	Up

6. Edit port3 (external2) interface.
7. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	172.22.222.4/255.255.255.0
Administrative Access	HTTPS SSH

Description	BGP external2 Peer2
Administrative Status	Up

To configure the FortiGate interfaces - CLI

```
config system interface
   edit port1
      set alias dmz
      set ip 10.11.201.110 255.255.255.0
      set allowaccess https ssh ping
      set description "OSPF internal networks"
      set status up
   next
   edit port2
      set alias external1
      set ip 172.21.111.5 255.255.255.0
      set allowaccess https ssh
      set description "external1 Peer 1"
      set status up
   next
   edit port3
      set alias external2
      set ip 172.22.222.5 255.255.255.0
      set allowaccess https ssh
      set description "external2 Peer 2"
      set status up
   next
end
```

To configure the firewall addresses - web-based manager

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New** and set the following information.

Category	Address
Name	BGP_services
Type	Subnet / IP Range
Subnet / IP Range	10.11.201.0 255.255.255.0
Interface	port1

3. Select **OK**.

To configure the firewall addresses - CLI

```
config firewall address
   edit "BGP_services"
      set associated-interface "port1"
      set subnet 10.11.201.0 255.255.255.0
   next
```

```
end
```

To configure firewall service groups - web-based manager

1. Go to **Policy & Objects > Objects > Services.** Under the **Create New** dropdown menu, select **Service Group**.
2. Name the group BGP_Services.
3. Add the following services to the **Members** list: BGP, DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
4. Select **OK**.
5. Create another new **Service Group**.
6. Name the group OSPF_Services.
7. Add the following services to the *Members* list: DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, OSPF, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
8. Select **OK**.

To configure firewall service groups - CLI

```
config firewall service group
   edit "BGP_services"
      set member "BGP", "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS" "IMAP"
          "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH" "TRACEROUTE" "SYSLOG"
   next
   edit "OSPF_services"
      set member "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS" "IMAP" "MYSQL"
          "NTP" "PING" "POP3" "SMTP" "SSH" "TRACEROUTE" "SYSLOG" "OSPF"
   next
end
```

Configuring BGP on the FortiGate unit

The only change from the standard BGP configuration for this example is configuring the blocking Peer 1's routes from being advertised to Peer 2. From the network topology you can guess that both of these peers likely share many routes in common and it does not make sense to advertise unneeded routes.

Blocking Peer 1's routes to Peer 2 is done with the distribute-list-out keyword. They allow you to select which routes you will advertise to a neighbor using an access list. In this case, we will block all incoming routes from Peer 1 when we send updates to Peer 2. Otherwise Peer 1 and Peer 2 are regular neighbors.

The FortiGate unit will redistribute routes learned from OSPF into BGP.

This is advanced configuration and the commands are only available in the CLI.

To create access list to block Peer 1 - CLI

```
config access-list
   edit "block_peer1"
      config rule
      edit 1
         set prefix 172.21.111.0 255.255.255.0
         set action deny
         set exact-match enable
      end
   end
```

```
end
```

To configure BGP on the FortiGate unit - CLI

```
config router bgp
   set as 65001
      set router-id 10.11.201.110
      config redistribute ospf
         set status enable
      end
      config neighbor
      edit 172.22.222.5
         set remote-as 65001
         set distribute-list-out "block_peer1"
      next
      edit 172.21.111.5
         set remote-as 65001
      end
   end
```

Configuring OSPF on the FortiGate unit

This configuration involves only one OSPF area, so all traffic will be intra-area. If there were two or more areas with traffic going between them, it would be inter-area traffic. These two types are comparable to BGP's traffic within one AS (iBGP) or between multiple ASes (eBPG). Redistributing routes from OSPF to BGP is considered external because either the start or end point is a different routing protocol.

The OSPF configuration is basic, apart from redistributing BGP routes learned.

To configure OSPF on the FortiGate unit - web-based manager

1. Go to **Router > Dynamic > OSPF**.
2. For Router ID enter `10.11.201.110` and then select **Apply**.
3. Under **Advanced Options** > **Redistribute**, select **BGP** and set the BGP **Metric** to 1.
4. For **Areas**, select **Create New,** enter the following information and then select **OK**.

Area (IP)	`0.0.0.0`
Type	Regular
Authentication	None

5. For **Networks**, select **Create New**.
6. Enter 10.11.201.0/255.255.255.0 for **IP/Netmask**, and select **OK**.
7. For **Interfaces**, select **Create New**.
8. Enter `OSPF_dmz_network` for **Name**.
9. Select `port1(dmz)` for **Interface** and then select **OK**.

To configure OSPF on the FortiGate unit - CLI

```
config router ospf
   set router-id 10.11.201.110
   config area
```

```
      edit 0.0.0.0
         set type regular
         set authentication none
      end
   config network
      edit 1
         set area 0.0.0.0
         set prefix 10.11.201.0 255.255.255.0
      end
   config interface
      edit "OSPF_dmz_network"
         set interface port1(dmz)
         set status enable
      end
   config redistribute bgp
      set status enable
      set metric 1
   end
end
```

Configuring other networking devices

As with all BGP configurations, the peer routers will need to be updated with the FortiGate unit's BGP information, including IP address, AS number, and what capabilities are being used, such as IPv6, graceful restart, BFD, and so on.

Configuring ECMP support for BGP

Equal Cost Multiple Path (ECMP) is a mechanism that allows multiple routes to the same destination with different next-hops and load-balances routed traffic over those multiple next-hops.

- ECMP only works for routes that are sourced by the same routing protocol (Static Routes, OSPF, and BGP).
- ECMP is enabled, by default, with 10 paths.
- ECMP with static routes is effective if the routes are configured with the same distance and same priority.

To configure ECMP support - CLI

```
config router bgp
   set ebgp-multipath disable[|enable]
   set ibgp-multipath disable[|enable]
   ...
end
```

Testing network configuration

Testing this configuration involves the standard connectivity checks, but also ensures that routes are being passed between protocols as expected.

Check the routing table on the FortiGate unit to ensure that routes from both OSPF and BGP are present.

Check the routing table on devices on the OSPF network for routes redistributed from BGP. Also, check those devices for connectivity to the Internet.

Check the routing table on Peer 2 to ensure that no routes from Peer 1 are present, but routes from the internal OSPF network are present.

For help with troubleshooting, see Redistributing and blocking routes in BGP on page 230.

IS-IS

This section describes the Intermediate System to Intermediate System Protocol (IS-IS).

IS-IS background and concepts

Intermediate System to Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) that is not intended to be used between Autonomous Systems (AS).

Background

IS-IS was developed by Digital Equipment Corporation and later standardized by ISO in 1992 as ISO 19589 (see RFC 1142, note that this RFC is different from the ISO version). About the same time, the Internet Engineering Task Force developed OSPF (see OSPF on page 158). After the initial version, IP support was added to IS-IS and this version was called Integrated IS-IS (see RFC 1195). Its widespread use started when an early version of IS-IS was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by IS-IS, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

IS-IS is a link state protocol that is well-suited to smaller networks. It is in widespread use and has near universal support on routing hardware. It is quick to configure and works well if there are no redundant paths. However, IS-IS updates are sent out node-by-node, so it can be slow to find a path around network outages. IS-IS also lacks good authentication, cannot choose routes based on different quality of service methods, and can create network loops if you are not careful. IS-IS uses Djikstra's algorithm to find the best path, like OSPF.

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures, largely due to its native support for IPv6 and its non-disruptive methods for splitting, merging, migrating, and renumbering network areas.

The FortiGate implementation supports IS-IS for IPv4 (see RFCs 1142 and 1162), but does not support IS-IS for IPv6 (although this technically can be achieved using the ZebOS routing module).

How IS-IS works

As one of the original modern dynamic routing protocols, IS-IS is straightforward. Its routing algorithm is not complex, there are some options to allow fine-tuning, and it is straightforward to configure IS-IS on FortiGate units.

From RFC 1142:

> *The routing algorithm used by the Decision Process is a shortest path first (SPF) algorithm. Instances of the algorithm are run independently and concurrently by all intermediate systems in a routing domain. IntraDomain routing of a PDU occurs on a hop-by-hop basis: that is, the algorithm determines only the next hop, not the complete path, that a data PDU will take to reach its destination.*

IS-IS versus static routing

IS-IS was one of the earliest dynamic routing protocols to work with IP addresses. As such, it is not as complex as more recent protocols. However, IS-IS is a big step forward from simple static routing.

While IS-IS may be slow in response to network outages, static routing has zero response. The same is true for convergence: static routing has zero convergence. Both IS-IS and static routing have the limited hop count, so it is neither a strength or a weakness.

TLV

IS-IS uses *type-length-variable (TLV)* parameters to carry information in Link-State PDUs (LSPs). Each IS-IS LSP consists of a variable-length header to which TLVs are appended in order to extend IS-IS for IP routing. The TLV field consists of one octet of type (T), one octet of length (L), and “L” octets of value (V). They are included in all of the IS-IS Packet types. For a complete breakdown of the LSP, see LSP structure on page 239.

In IS-IS, TLVs are used to determine route-leaking and authentication and are also used for IPv4 and IPv6 awareness and reachability.

- To determine which TLVs are responsible for route-leaking, see Default routing on page 241.
- To determine which TLVs are responsible for authentication, see Authentication on page 243.

For a complete list of reserved TLV codepoints, refer to RFC 3359.

LSP structure

It is difficult to fully understand a routing protocol without knowing what information is carried in its packets. Knowing how routers exchange each type of information will help you better understand the IS-IS protocol and will allow you to configure your network more appropriately.

This section provides information about the contents of the IS-IS LSP. LSPs describe the network topology and can include IP routes and checksums.

NSAP and NET

IS-IS routing protocol utilizes ISO network addressing to identify network interfaces. The addresses are known as Network Service Access Points (NSAP). In general, IS-IS routers consist of only one NSAP, whereas IP addressing requires one IP address per interface.

In IS-IS, the NSAP address is translated into a Network Entity Title (NET), which is the same as the NSAP but can differentiate end systems by way of a byte called the *n-selector* (NSEL). In order for adjacencies to form in IS-IS, the NSEL must be set to zero, to indicate “this system”. The total NET can be anywhere between 8 and 20 bytes long due to the support for variable length area addressing.

The following diagram identifies the individual parts of the NSAP, with explanations below:

NSAP and NET example

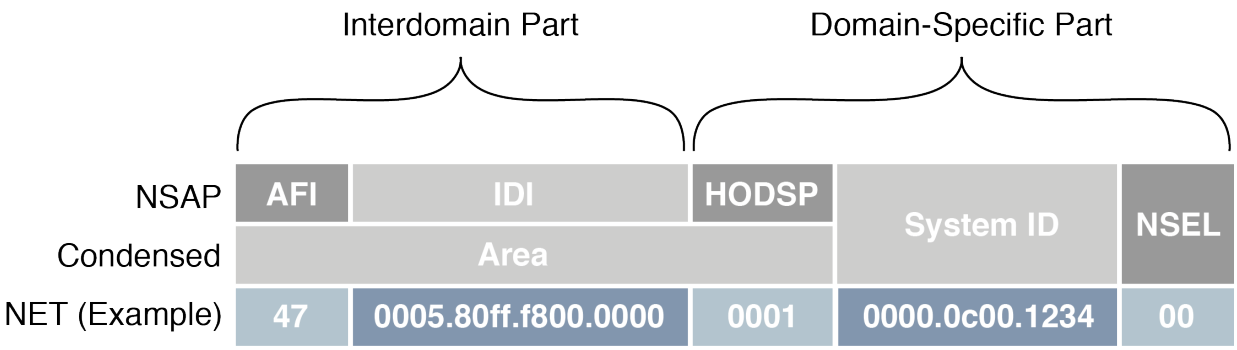

- **AFI** : The *Authority and Format Identifier (AFI)* specifies the format of the addressing family used. IS-IS is designed to carry routing information for several different protocols. Each entry has an address family identifier that identifies the globally unique Interdomain Part (IDP). For example, 49 is the AFI for private addresses, whereas 47 is the AFI for international organizations.
- **IDI**: The *Initial Domain Identifier (IDI)* identifies the routing domain within an interconnected network. The length of the IDI is typically determined by the AFI. If you are using an AFI of 49, you do not need to specify an IDI since the network is private.
- **HODSP** : The *High Order Domain-Specific Part (HODSP)* identifies the unique address within a specific routing domain. Together, the AFI, IDI, and HODSP define the area address. All of the nodes within an area must have the same area address.
- **System ID** : The *System ID* represents the 6-8 byte router identifier. The ID could be Media Access Control (MAC) format, as in the example above, or a static length IP address expressed in binary-coded decimal (BCD) format.
- **NSEL**: The *n-selector (NSEL)*, as previously described, identifies the network layer transport service and must always be set to zero for IS-IS NETs.

Parts and terminology of IS-IS

Before you can understand how IS-IS functions, you need to understand some of the main concepts and parts of IS-IS.

DIS election and pseudonode LSP

In IS-IS routing protocol, a single router is chosen to be the designated intermediate system (DIS). The election of the DIS is determined automatically and dynamically on the LAN depending on highest interface priority and the subnetwork point of attachment (SNPA). The FortiGate is typically the DIS, and each router in its LAN is an intermediate system (IS).

Unlike OSPF, which elects a designated router (DR) and backup designated router (BDR), the DIS has no backup and determines the election of a new DIS whenever a router is added to the LAN or whenever the current DIS drops. A backup DIS is irrelevant since all of the routers on an IS-IS system are synchronized, and the short Hello interval used by the DIS quickly detects failures and the subsequent replacement of the DIS.

Synchronization of all the nodes in an IS-IS area could prove troublesome when updating the network infrastructure and would demand ever-increasing resources each time a new router is added (at an exponential scale). For this purpose, the DIS creates a pseudonode, which is essentially a virtual, logical node representing the LAN. The pseudonode requests adjacency status from all the routers in a multi-access network by sending IS-

IS Hello (IIH) PDUs to Level 1 and Level 2 routers (where Level 1 routers share the same address as the DIS and Level 2 routers do not). Using a pseudonode to alter the representation of the LAN in the link-state database (LSD) greatly reduces the amount of adjacencies that area routers have to report. In essence, a pseudonode *collapses* a LAN topology, which allows a more linear scale to link-state advertising.

In order to maintain the database synchronization, the DIS periodically sends complete sequence number packets (CSNPs) to all participating routers.

Packet types

Four general packet types (PDUs) are communicated through IS-IS, appearing at both Level 1 and Level 2. They are described below.

- **Intermediate System-to-Intermediate System Hello (IIH) PDU** : As mentioned previously, the IIH PDU, or Hello packet, detects neighboring routers and indicates to the pseudonode the area's adjacency mesh. The Hello packet, flooded to the multicast address, contains the system ID of the sending router, the holding time, the circuit type of the interface on which the PDU was sent, the PDU length, the DIS identifier, and the interface priority (used in DIS election). The Hello packet also informs its area routers that it is the DIS. Hello packets are padded to the maximum IS-IS PDU size of 1492 bytes (the full MTU size) to assist in the detection of transmission errors with large frames or with MTU mismatches between adjacencies. The DIS typically floods Hello packets to the entire LAN every three seconds.
- **Link-state PDU (LSP)** : The LSP contains information about each router in an area and its connected interfaces. LSPs are refreshed periodically and acknowledged on the network by way of sequence number PDUs. If new LSP information is found, based on the most recent complete sequence number PDU (CSNP), out-of-date entries in the link-state database (LSDB) are removed and the LSDB is updated. For a more detailed breakdown of the LSP, see LSP structure on page 239.
- **Complete sequence number PDU (CSNP)**: CSNPs contain a list of all LSPs in the current LSDB. The CSNP informs other area routers of missing or outdated links in the adjacency mesh. The receiving routers then use this information to update their own database to ensure that all area routers converge. In contrast to Hello packets, CSNPs are sent every ten seconds and only between neighbors. In other words, they are never flooded.
- **Partial sequence number PDU (PSNP)** : PSNPs are used to request and acknowledge LSP information from an adjacency. When a router compares a CSNP with its local database and determines a discrepancy, the router requests an updated LSP using a PSNP. Once received, the router stores the LSP in its local database and responds to the DIS with acknowledgement.

Default routing

The default route is used if there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

By default, all routes are displayed in the Routing Monitor list. To display the routes in the routing table, go to **Monitor > Routing Monitor**.

Route leaking

Route leaking is the term used to describe the bi-directional flow of information between internal and external routing interfaces. By default, IS-IS leaks routing information from a Level 1 area into a Level 2 area. In order to

leak Level 2 routing information into a Level 1 area, you must configure an export policy. Whether or not a route is leaked is determined by the ATT bit, using TLV 128 (for internal IP reachability) and TLV 130 (for external IP address information). For more information about TLVs, see Troubleshooting IS-IS on page 244.

To configure IS-IS route leaking, use the following CLI commands.

1. On a Level 1-2 router:
```
config router isis
   set redistribute-l2 enable
end
```
2. On a Level 1 router:
```
config router isis
   get router info routing-table isis
   get router info isis route
end
```

Default information originate option

Enabling default-information-originate generates and advertises a default route into the FortiGate unit's IS-IS-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. IS-IS does not create the default route unless you use the `always` option.

Select **Disable** if you experience any issues or if you wish to advertise your own static routes into IS-IS updates.

The CLI commands associated with default information originate include:

```
config router isis
   set default-originate
end
```

Timer options

IS-IS uses various timers to regulate its performance, including garbage, update, and timeout timers. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations. If you change these settings, ensure that the new settings are compatible with local routers and access servers.

You can configure the three IS-IS timers in the CLI, using the following commands:

```
config router isis
   set garbage-timer
   set update-timer
   set timeout-timer
end
```

You will find more information on each timer below.

Update timer

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There is some randomness added to help prevent network traffic congestion, which could result from all routers simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, or you will experience an error.

If you are experiencing significant traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you will experience timeouts that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the DIS will keep a reachable route in the routing table while no updates for that route are received. If the DIS receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period, or you will experience an error.

If you are experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods. It may be a considerable amount of time before the DIS is done waiting for all the timers to expire on unresponsive routes.

Garbage timer

The garbage timer is the amount of time (in seconds) that the DIS will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This results in a smaller routing table, which is useful if you have a very large network or if your network changes frequently.

Authentication

In routing protocols, it is typically desirable to establish authentication rules that prevent malicious and otherwise unwanted information from being injected into the routing table. IS-IS routing protocol utilizes TLV 10 to establish authentication. For more information about TLVs, see TLV on page 239.

Initially, IS-IS used plain cleartext to navigate the authentication rules, but this was found to be insecure since the cleartext packets were unencrypted and could be exposed to packet sniffers. As per RFC 3567, HMAC-MD5 and enhanced cleartext authentication features were introduced in IS-IS, both of which encrypt authentication data, making them considerably more secure than using plain cleartext authentication.

HMAC-MD5 authentication

Hashed Message Authentication Codes - Message Digest 5 (HMAC-MD5) is a mechanism for applying a cryptographic hash function to the message authentication process. It is applied at both Level 1 and Level 2 routing. In IS-IS, an HMAC-MD5 can be applied to each type of LSP, on different interfaces, and with different passwords.

Authentication data is hashed using an AH (Authentication Header) key. From RFC 2085:

> *The "AH Key" is used as a shared secret between two communicating parties. The Key is not a "cryptographic key" as used in a traditional sense. Instead, the AH key (shared secret) is hashed with the transmitted data and thus, assures that an intervening party cannot duplicate the authentication data. [...] Implementation should, and as frequently as possible, change the AH key. Keys need to be chosen at random, or generated using a cryptographically strong pseudo-random generator seeded with a random seed."*

Cleartext authentication uses the configuration commands `area-password` and `domain-password` for authentication, but when migrating from cleartext authentication to HMAC-MD5, these command settings are automatically overwritten.

By the year 2005, the MD5 hash function had been identified as vulnerable to collision search attacks and various weaknesses. While such vulnerabilities do not compromise the use of MD5 within HMAC, administrators need to be aware of potential developments in cryptanalysis and cryptographic hash functions in the likely event that the underlying hash function needs to be replaced.

Enhanced cleartext authentication

Enhanced cleartext authentication is an extension to cleartext authentication that allows the encryption of passwords as they are displayed in the configuration. It includes a series of authentication mode commands and an authentican key chain, and allows for more simple password modification and password management. Enhanced cleartext authentication also provides for smoother migration to and from changing authentication types. Intermediate systems continue to use the original authentication method until all the area routers are updated to use the new method.

Authentication key chain

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. A router migrates from one key to the next according to the scheduled send and receive lifetimes. If an active key is unavailable, the PDU is automatically discarded.

From RFC 5310:

> *It should be noted that the cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function and on the size and quality of the key.*

Troubleshooting IS-IS

Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there is a routing loop, that normal path doubles back on itself creating a loop. When there are loops, the network has problems.

A routing loop happens when a normally functioning network has an outage and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it is possible for a route to be attempted that involves going back a hop and trying a different hop forward. If that hop forward is blocked by the outage as well, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on the affected routers. The worst part is this situation will continue until the network administrator changes the router settings or the downed routers come back online.

Routing loop effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes that router sends an update out to all of the IS-IS routers connected to it. In a network loop, it is possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How to spot a routing loop

Any time network traffic slows down, you will ask yourself if it is a network loop or not. Often slowdowns are normal. They are not a full stoppage and normal traffic resumes in a short period of time.

If the slowdown is a full halt of traffic, or a major slowdown does not return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- Checking your logs
- Using SNMP network monitoring
- Using Link Health Monitor and e-mail alerts
- Looking at the packet flow

If you are not running SNMP or dead gateway detection, or if you have non-Fortinet routers in your network, you can use networking tools such as ping and traceroute to define the outage on your network and begin to fix it.

Checking your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On your FortiGate unit, go to **Log & Report > Log & Archive Access**. You will want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to dead gateway detection), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensic analysis can better help you prepare for next time.

Using SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it is exactly, as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops

1. Go to **System > Config > SNMP**.
2. Enable **SNMP Agent**.
3. Optionally enter the **Description**, **Location**, and **Contact** information for this device for easier location of the problem report.
4. In either **SNMP v1/v2c** section or **SNMP v3** section, as appropriate, select **Create New**.
5. Enter the **Community Name** that you want to use.
6. In **Hosts**, select **Add** to add an IP address where you will be monitoring the FortiGate unit. You can add up to 8 different addresses.
7. Ensure that ports 161 and 162 (SNMP queries and traps) are allowed through your security policies.

8. In **SNMP Event**, select the events you want to be notified about. For routing loops, this should include **CPU Overusage**, **Memory Low**, and possibly **Log disk space low**. If there are problems, the log will fill up quickly, and the FortiGate unit's resources will be overused.
9. Select **OK**.
10. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate unit. Typically, you can configure this software to alert you about outages or CPU spikes that may indicate a routing loop.

Using Link Health Monitor and e-mail alerts

Another tool available to you on FortiGate units is the Link Health Monitor. This is useful for dead gateway detection. This feature allows the FortiGate unit to ping a gateway at regular intervals to ensure it is online and working. When the gateway is not accessible, that interface is marked as down.

To detect possible routing loops with Link Health Monitor

Use the following command to configure dead gateway detection:

```
config system link-monitor
   edit "test"
      set srcintf "internal4"
      set server "8.8.8.8"
      set interval 5
      set failtime 1
   end
```

Set the `Interval` (how often to send a ping) and `failtime` (how many lost pings is considered a failure). A smaller interval and smaller number of lost pings will result in faster detection, but will create more traffic on your network.

You may also want to log CPU and memory usage, as a network outage will cause your CPU activity to spike.

If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate unit cannot connect to the next router, the FortiGate unit will bring down the interface and alert you with an email about the outage.

Looking at the packet flow

If you want to see what is happening on your network, look at the packets travelling on the network. In this situation, you are looking for routes that have metrics higher than 15 as that indicates they are unreachable. Ideally if you debug the flow of the packets, and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you have mapped the problem on your network and determined it is in fact a routing loop, there are a number of steps to take to correct it.

1. Get any offline routers back online. This may be a simple reboot or you may have to replace hardware. Often this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Split horizon and poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let us call them routerA, routerB, and routerC. RouterA is linked only to routerB, routerC is linked only to routerB, and routerB is linked to both routerA and routerC. To get to routerC, routerA must go through routerB. If the link to routerC goes down, it is possible that routerB will try to use routerA's route to get to routerC. This route is A-B-C, so it will not work. However, if routerB tries to use it, this begins an endless loop. This situation is called a split horizon because from routerB's point of view, the horizon stretches out in each direction but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This "poisoned" route is marked as unreachable for routers that cannot use it. In IS-IS, this means that route is marked with a distance of 16.

Simple IS-IS example

This is an example of a typical medium-sized network configuration using IS-IS routing.

Imagine a company with four FortiGate devices connected to one another. A FortiGate at one end of the network connects to two routers, each with its own local subnet. One of these routers uses OSPF and the other router uses RIP.

Your task is to configure the four FortiGates to route traffic and process network updates using IS-IS, so that the farthest FortiGate (see 'FGT4' in Network layout and assumptions on page 248) receives route updates for the two routers at the opposite end of the network. Furthermore, FGT4 has been given a loopback subnet that must be identified by the router running RIP.

Since the internal networks use OSPF and RIP, those protocols will need to be redistributed through the IS-IS network. To keep the example simple, there will be no authentication of router traffic.

With IS-IS properly configured in this example, if a router fails or temporarily goes offline, the route change will propagate throughout the system.

Network layout and assumptions

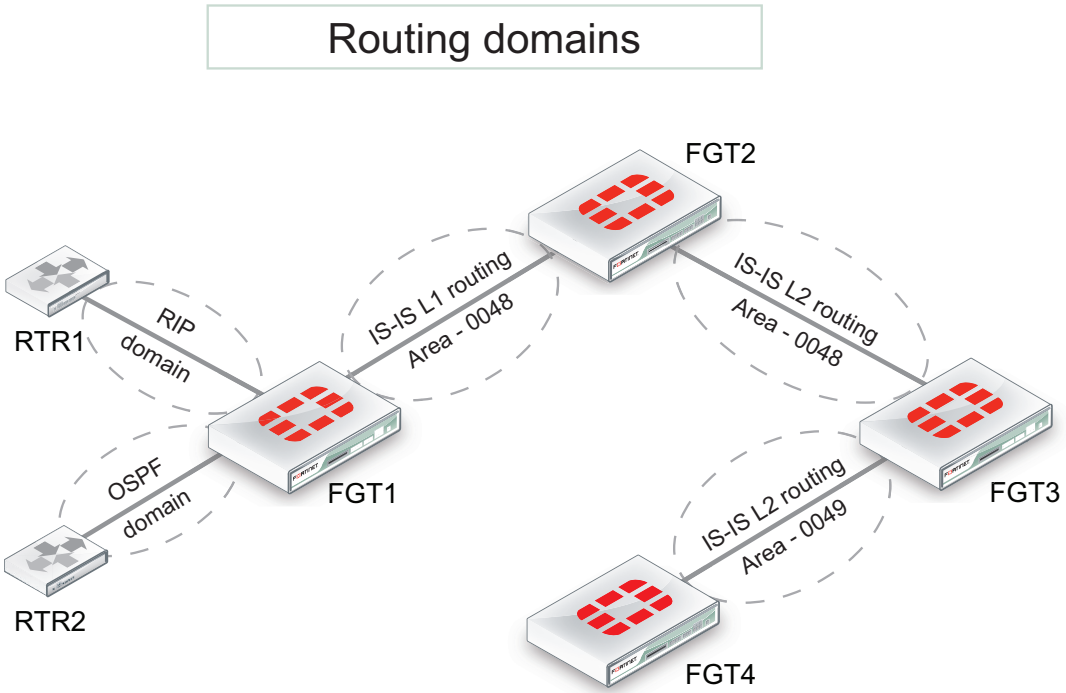

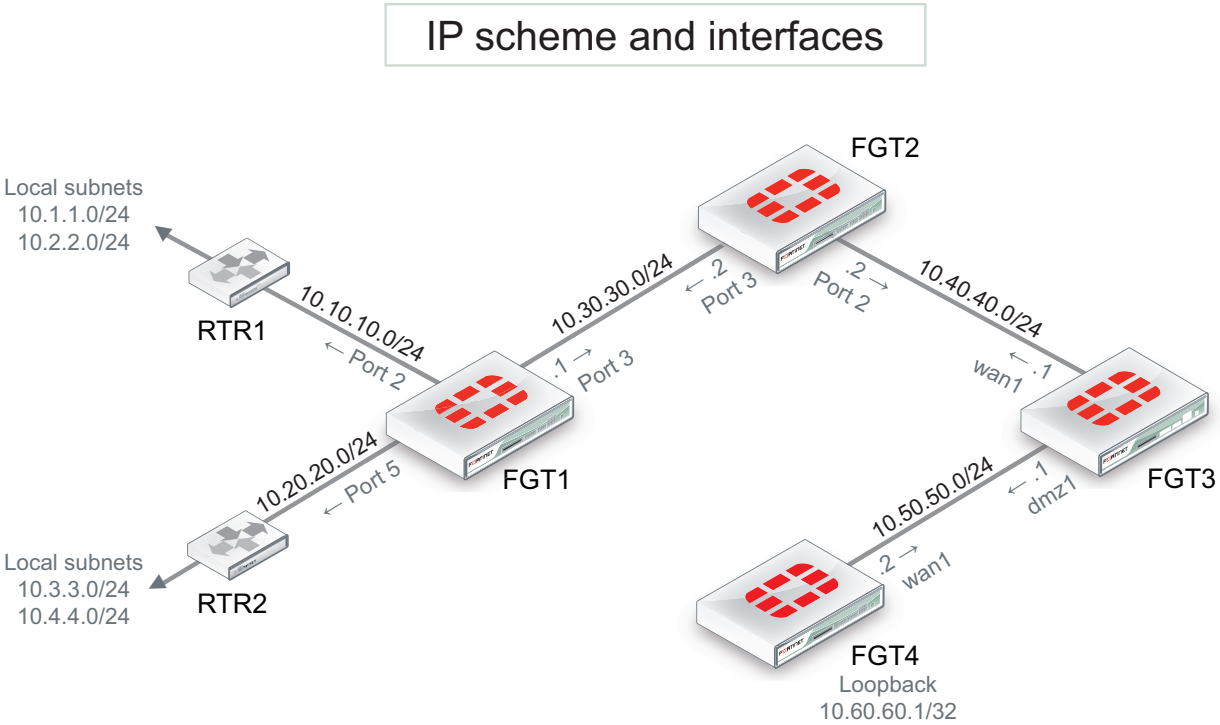

- It is assumed that each FortiGate is operating in NAT mode, running FortiOS 4.0MR2+.
- All interfaces have been previously assigned and no static routes are required.

- The Authority and Format Identifier (AFI) used is 49 : Locally administered (private).
- The Area identifiers are 0048 and 0049.

Expectations

- FGT4 must get the IS-IS route updates for RTR1 and RTR2 local subnets (10.1.1.0, 10.2.2.0, 10.3.3.0, 10.4.4.0).
- RTR1 must receive (via RIP2) the loopback subnet of FGT4 (10.60.60.1/32).

CLI configuration

The following CLI configuration occurs on each FortiGate (as identified), including only the relevant parts.

FGT1

```
config router isis
   config isis-interface
      edit "port3"
         set circuit-type level-1
         set network-type broadcast
         set status enable
      next
   end
   config isis-net
      edit 1
         set net 49.0048.1921.6818.2136.00
      next
   end
   config redistribute "connected"
   end
   config redistribute "rip"
      set status enable
      set level level-1
   end
   config redistribute "ospf"
      set status enable
      set level level-1
   end
end
config router rip
   config interface
      edit "port2"
         set receive-version 2
         set send-version 2
      next
   end
   config network
      edit 1
         set prefix 10.10.10.0 255.255.255.0
      next
   end
   config redistribute "isis"
      set status enable
   end
end
```

FGT2

```
config router isis
   config isis-interface
      edit "port3"
         set circuit-type level-1
         set network-type broadcast
         set status enable
      next
      edit "port2"
         set network-type broadcast
         set status enable
      next
   end
   config isis-net
      edit 1
         set net 49.0048.1221.6818.2110.00
      next
   end
   set redistribute-l1 enable
   set redistribute-l2 enable
end
```

FGT3

```
config router isis
   set is-type level-2-only
      config isis-interface
         edit "wan1"
            set network-type broadcast
            set status enable
         next
         edit "dmz1"
            set network-type broadcast
            set status enable
         next
      end
      config isis-net
         edit 1
            set net 49.0048.1921.6818.2108.00
         next
         edit 2
            set net 49.0049.1921.6818.2108.00
         next
      end
end
```

FGT4

```
config router isis
   set is-type level-2-only
      config isis-interface
         edit "wan1"
            set network-type broadcast
            set status enable
```

```
            next
         end
         config isis-net
            edit 1
               set net 49.0049.1721.0160.1004.00
            next
         end
         config redistribute "connected"
            set status enable
         end
   end
```

Verification

Once the network has been configured, you need to test that it works as expected. Use the following CLI commands on the devices indicated.

Verifying if RTR1 receives loopback subnet of FGT4

```
(RTR1) # get router info routing-table all
```

Result:

```
C    10.1.1.0/24 is directly connected, vlan1
C    10.2.2.0/24 is directly connected, vlan2
C    10.10.10.0/24 is directly connected, dmz1
R    10.40.40.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R    10.50.50.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R    10.60.60.1/32 [120/2] via 10.10.10.1, dmz1, 00:04:07
```

(*) If required, filtering out 10.50.50.0 and 10.40.40.0 from the routing table could be done with a route-map.

Verification on FGT2, which is the border between L1 and L2 routing levels; looking at IS-IS information

```
FGT2 # get router info isis interface
```

Result:

```
port2 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x01
Extended Local circuit ID: 0x00000003
Local SNPA: 0009.0f85.ad8c
IP interface address:
10.40.40.2/24
IPv4 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 6 seconds
Next IS-IS LAN Level-2 Hello in 1 seconds
port3 is up, line protocol is up
```

```
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000004
Local SNPA: 0009.0f85.ad8d
IP interface address:
10.30.30.2/24
IPv4 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.02
Number of active level-1 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 2 seconds

FGT2 # get router info isis neighbor
```

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6818.2108	port2	0009.0f04.0794	Up	22	L2	IS-IS
1921.6818.2136	port3	0009.0f85.acf7	Up	29	L1	IS-IS

Verification on FGT3, which is border between 2 areas, looking at IS-IS information

IS-IS router CLI commands available:

```
FGT3 # get router info isis ?
```

Result:

```
interface                        show isis interfaces

neighbour                        show CLNS neighbor adjacencies

is-neighbour                     show IS neighbor adjacencies

database                         show IS-IS link state database

route                            show IS-IS IP routing table

topology                         show IS-IS paths
```

Example of interface status and neighbors:

```
FGT3 # get router info isis interface
```

Result:

```
wan1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
```

```
Local circuit ID: 0x01
Extended Local circuit ID: 0x00000003
Local SNPA: 0009.0f04.0794
IP interface address:
10.40.40.1/24
IPv4 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 3 seconds

dmz1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000005
Local SNPA: 0009.0f04.0792
IP interface address:
10.50.50.1/24
IPv4 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1721.0160.1004.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 7 seconds

FGT3 # get router info isis neighbor
```

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1221.6818.2110	wan1	0009.0f85.ad8c	Up	8	L2	IS-IS
1721.0160.1004	dmz1	0009.0f52.7704	Up	8	L2	IS-IS

Verification on FGT4 that the remote subnets from RTR1 and RTR2 are in the routing table and learned with IS-IS

```
FGT4 # get router info routing-table all
```

Result:

```
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
     N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
     E1 - OSPF external type 1, E2 - OSPF external type 2
     i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
      * - candidate default
i L2 10.1.1.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
    i L2 10.2.2.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
    i L2 10.3.3.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
    i L2 10.4.4.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.10.10.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
    i L2 10.11.11.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
    i L2 10.20.20.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
```

```
    i L2 10.30.30.0/24 [115/30] via 10.50.50.1, wan1, 00:13:55
    i L2 10.40.40.0/24 [115/20] via 10.50.50.1, wan1, 00:15:30
C 10.50.50.0/24 is directly connected, wan1
    C 10.60.60.1/32 is directly connected, loopback
```

Troubleshooting

The following diagnose commands are available for further IS-IS troubleshooting and will display all IS-IS activity (sent and received packets):

```
FGT # diagnose ip router isis level info
FGT # diagnose ip router isis all enable
FGT # diagnose debug enable
```

…to stop the debug type output:

```
FGT # diagnose ip router isis level none
```

Output and interpretation depends on the issue faced. You can provide this information to TAC if you open a support ticket.

Multicast forwarding

Multicasting (also called IP multicasting) consists of using a single multicast source to send data to many receivers. Multicasting can be used to send data to many receivers simultaneously while conserving bandwidth and reducing network traffic. Multicasting can be used for one-way delivery of media streams to multiple receivers and for one-way data transmission for news feeds, financial information, and so on.

Also, RIPv2 uses multicasting to share routing table information, OSPF uses multicasting to send hello packets and routing updates, Enhanced Interior Gateway Routing Protocol (EIGRP) uses multicasting to send routing information to all EIGRP routers on a network segment and the Bonjour network service uses multicasting for DNS.

A FortiGate unit can operate as a Protocol Independent Multicast (PIM) version 2 router. FortiGate units support PIM sparse mode (RFC 4601) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate unit interface is connected. Multicast routing is not supported in transparent mode (TP mode).

To support PIM communications, the sending/receiving applications and all connecting PIM routers in between must be enabled with PIM version 2. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations. To enable source-to-destination packet delivery, either sparse mode or dense mode must be enabled on the PIM-router interfaces. Sparse mode routers cannot send multicast messages to dense mode routers. In addition, if a FortiGate unit is located between a source and a PIM router, two PIM routers, or is connected directly to a receiver, you must create a security policy manually to pass encapsulated (multicast) packets or decapsulated data (IP traffic) between the source and destination.

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Boot Strap Router (BSR), and if sparse mode is enabled, a number of Rendezvous Points (RPs) and Designated Routers (DRs). When PIM is enabled on a FortiGate unit, the FortiGate unit can perform any of these functions at any time as configured.

Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast group(s) that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.

When a FortiGate unit interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP's distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that is associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with RFC 4601: Protocol Independent Multicast - Sparse Mode (PIM-SM). You can use the following CLI commands to configure IPv6 PIM sparse multicast routing:

```
config router multicast6
   set multicast-routing {enable | disable}
      config interface
         edit <interface-name>
            set hello-interval <1-65535 seconds>
            set hello-holdtime <1-65535 seconds>
         end
      config pim-sm-global
         config rp-address
            edit <index>
               set ipv6-address <ipv6-address>
            end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information, if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly connected network segments. The PIM routers exchange state refresh messages to update their distribution trees. FortiGate units store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that do not have downstream receivers. PIM routers that do not manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

FortiGate units operating in NAT mode can also be configured as multicast routers. You can configure a FortiGate unit to be a Protocol Independent Multicast (PIM) router operating in Sparse Mode (SM) or Dense Mode (DM).

PIM support

A FortiGate unit can be configured to support PIM by going to **Network > Multicast** and enabling multicast routing. You can also enable multicast routing using the `config router multicast` CLI command. When PIM is enabled, the FortiGate unit allocates memory to manage mapping information. The FortiGate unit communicates with neighboring PIM routers to acquire mapping information and if required, processes the multicast traffic associated with specific multicast groups.

The end-user multicast client-server applications must be installed and configured to initiate Internet connections and handle broadband content such as audio and video information.

Client applications send multicast data by registering IP traffic with a PIM-enabled router. An end user can type in a class D multicast group address, an alias for the multicast group address, or a conference call number to initiate the session.

Rather than sending multiple copies of generated IP traffic to more than one specific IP destination address, PIM-enabled routers encapsulate the data and use the one multicast group address to forward multicast packets to multiple destinations. Because one destination address is used, a single stream of data can be sent. Client applications receive multicast data by requesting that the traffic destined for a certain multicast group address be delivered to them. End users may use phone books, a menu of ongoing or future sessions, or some other method through a user interface to select the address of interest.

A class D address in the 224.0.0.0 to 239.255.255.255 range may be used as a multicast group address, subject to the rules assigned by the Internet Assigned Numbers Authority (IANA). All class D addresses must be assigned in advance. Because there is no way to determine in advance if a certain multicast group address is in use, collisions may occur (to resolve this problem, end-users may switch to a different multicast address).

To configure a PIM domain

1. If you will be using sparse mode, determine appropriate paths for multicast packets.
2. Make a note of the interfaces that will be PIM-enabled. These interfaces may run a unicast routing protocol.
3. If you will be using sparse mode and want multicast packets to be handled by specific (static) RPs, record the IP addresses of the PIM-enabled interfaces on those RPs.
4. Enable PIM version 2 on all participating routers between the source and receivers. On FortiGate units, use the `config router multicast` command to set global operating parameters.
5. Configure the PIM routers that have good connections throughout the PIM domain to be candidate BSRs.
6. If sparse mode is enabled, configure one or more of the PIM routers to be candidate RPs.
7. If required, adjust the default settings of PIM-enabled interface(s).

Multicast forwarding and FortiGate units

In both transparent mode and NAT mode, you can configure FortiGate units to forward multicast traffic.

For a FortiGate unit to forward multicast traffic, you must add FortiGate multicast security policies. Basic multicast security policies accept any multicast packets at one FortiGate interface and forward the packets out another FortiGate interface. You can also use multicast security policies to be selective about the multicast traffic that is accepted, based on source and destination address, and to perform NAT on multicast packets.

In the example shown below, a multicast source on the marketing network with IP address 192.168.5.18 sends multicast packets to the members of network 239.168.4.0. At the FortiGate unit, the source IP address for multicast packets originating from workstation 192.168.5.18 is translated to 192.168.18.10. In this example, the FortiGate unit is not acting as a multicast router.

Multicast forwarding and RIPv2

RIPv2 uses multicast to share routing table information. If your FortiGate unit is installed on a network that includes RIPv2 routers, you must configure the FortiGate unit to forward multicast packets so that RIPv2 devices can share routing data through the FortiGate unit. No special FortiGate configuration is required to share RIPv2 data, you can simply use the information in the following sections to configure the FortiGate unit to forward multicast packets.

RIPv1 uses broadcasting to share routing table information. To allow RIPv1 packets through a FortiGate unit, you can add standard security policies. Security policies to accept RIPv1 packets can use the ANY predefined firewall service or the RIP predefined firewall service.

Example multicast network including a FortiGate unit that forwards multicast packets

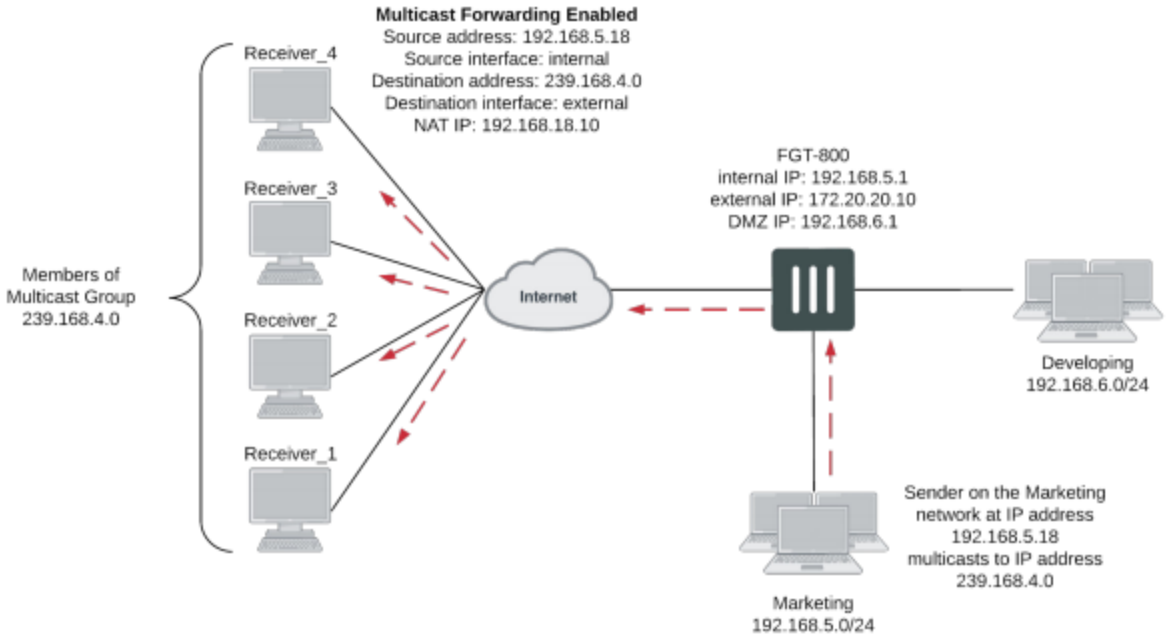

Configuring FortiGate multicast forwarding

You configure FortiGate multicast forwarding from the CLI. Two steps are required:

1. Adding multicast security policies
2. Enabling multicast forwarding
   This second step is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

There is sometimes confusion between the terms “forwarding” and “routing”. These two functions should not be taking place at the same time.

It is mentioned that multicast-forward should be enabled when the FortiGate unit is in NAT mode and that this will forward any multicast packet to all interfaces. However, this parameter should **NOT** be enabled when the FortiGate unit operates as a multicast router (for example, with a routing protocol enabled. It should only be enabled when there is no routing protocols activated.

Adding multicast security policies

You need to add security policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. You add multicast security policies from the CLI using the `config firewall multicast-policy` command. As with unicast security policies, you specify the source and destination interfaces and, optionally, the allowed address ranges for the source and destination addresses of the packets.

You can also use multicast security policies to configure source NAT and destination NAT for multicast packets.

Keep the following in mind when configuring multicast security policies:

- The matched forwarded (outgoing) IP multicast source IP address is changed to the configured IP address.
- Source and destination interfaces are optional. If left blank, the multicast will be forwarded to ALL interfaces.
- Source and destination addresses are optional. If left unset, it means ALL addresses.
- The `nat` keyword is optional. Use it when source address translation is needed.

Enabling multicast forwarding

Multicast forwarding is enabled by default. In NAT mode you must use the `multicast-forward` keyword of the `system settings` CLI command to enable or disable multicast forwarding. When `multicast-forward` is enabled, the FortiGate unit forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header will be reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add security policies to actually allow multicast packets through the FortiGate. In our example, the security policy allows multicast packets received by the internal interface to exit to the external interface.

Enabling multicast forwarding is only required if your FortiGate unit is operating in NAT mode. If your FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

Enter the following CLI command to enable multicast forwarding:

```
config system settings
   set multicast-forward enable
end
```

If multicast forwarding is disabled and the FortiGate unit drops packets that have multicast source or destination addresses.

You can also use the `multicast-ttl-notchange` keyword of the `system settings` command so that the FortiGate unit does not increase the TTL value for forwarded multicast packets. You should use this option only if packets are expiring before reaching the multicast router.

```
config system settings
   set multicast-ttl-notchange enable
end
```

In transparent mode, the FortiGate unit does not forward frames with multicast destination addresses. Multicast traffic, such as the one used by routing protocols or streaming media, may need to traverse the FortiGate unit and should not interfere with the communication. To avoid any issues during transmission, you can set up multicast security policies. These types of security policies can only be enabled using the CLI.

The CLI parameter multicast-skip-policy must be disabled when using multicast security policies. To disable enter the commands:

```
config system settings
     set multicast-skip-policy disable
end
```

In this simple example, no check is performed on the source or destination interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces on the forwarding domain, except the incoming interface.

To enable the multicast policy

```
config firewall multicast-policy
   edit 1
      set action accept
   end
```

In this example, the multicast policy only applies to the source port of WAN1 and the destination port of Internal.

To enable the restrictive multicast policy

```
config firewall multicast-policy
   edit 1
      set srcintf wan1
      set dstinf internal
      set action accept
   end
```

In this example, packets are allowed to flow from WAN1 to Internal, and sourced by the address 172.20.120.129, which is represented by the address object "example_addr-1".

To enable the restrictive multicast policy

```
config firewall multicast-policy
   edit 1
      set srcintf wan1
      set srcaddr example_addr-1
      set dstinf internal
      set action accept
   end
```

This example shows how to configure the multicast security policy required for the configuration shown. This policy accepts multicast packets that are sent from a PC with IP address 192.168.5.18 to destination address range 239.168.4.0. The policy allows the multicast packets to enter the internal interface and then exit the external interface. When the packets leave the external interface, their source address is translated to 192.168.18.10

```
config firewall multicast-policy
   edit 5
      set srcaddr 192.168.5.18 255.255.255.255
      set srcintf internal
      set destaddr 239.168.4.0 255.255.255.0
      set dstintf external
      set nat 192.168.18.10
   end
```

This example shows how to configure a multicast security policy so that the FortiGate unit forwards multicast packets from a multicast server with an IP 10.10.10.10 is broadcasting to address 225.1.1.1. This server is on the network connected to the FortiGate DMZ interface.

```
config firewall multicast-policy
   edit 1
      set srcintf DMZ
      set srcaddr 10.10.10.10 255.255.255.255
      set dstintf Internal
      set dstaddr 225.1.1.1 255.255.255.255
      set action accept
   edit 2
      set action deny
   end
```

Displaying IPv6 multicast router information

You can use the following CLI command to display IPv6 multicast router information (equivalent to the IPv4 version of the command):

```
get router info6 multicast
```

Multicast routing examples

This section contains the following multicast routing configuration examples and information:

- Example FortiGate PIM-SM configuration using a static RP
- FortiGate PIM-SM debugging examples
- Example multicast destination NAT (DNAT) configuration
- Example PIM configuration that uses BSR to find the RP

Example FortiGate PIM-SM configuration using a static RP

The example Protocol Independent Multicast Sparse Mode (PIM-SM) configuration shown below has been tested for multicast interoperability using PIM-SM between Cisco 3750 switches running 12.2 and a FortiGate-800 running FortiOS v3.0 MR5 patch 1. In this configuration, the receiver receives the multicast stream when it joins the group 233.254.200.1.

Example: FortiGate PIM-SM topology

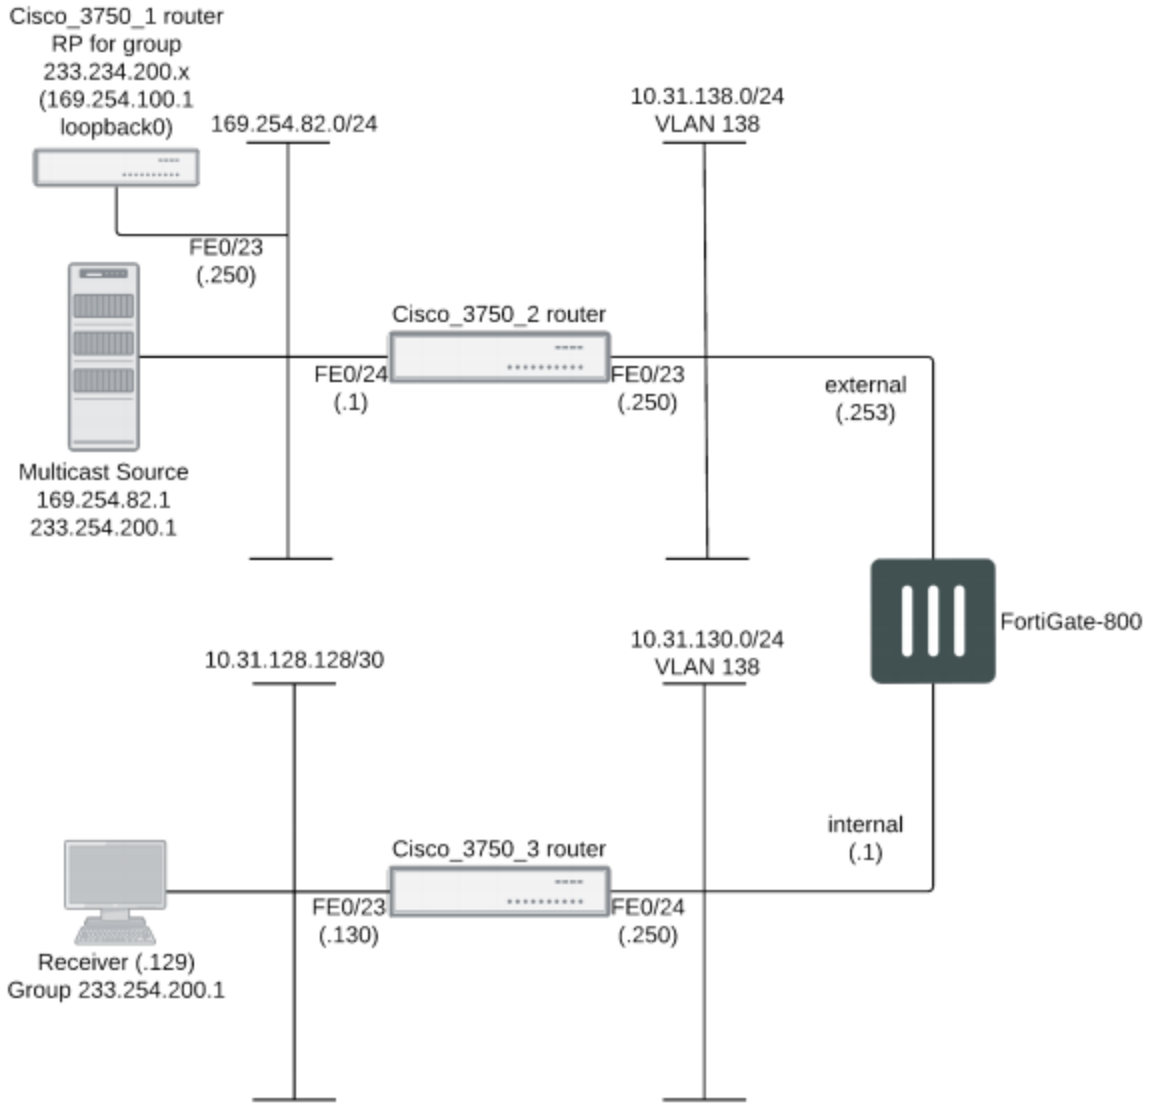

The configuration uses a statically configured rendezvous point (RP) which resides on the Cisco_3750_1. Using a bootstrap router (BSR) was not tested in this example. See “Example PIM configuration that uses BSR to find the RP” for an example that uses a BSR.

Configuration steps

The following procedures show how to configure the multicast configuration settings for the devices in the example configuration.

- Cisco_3750_1 router configuration
- Cisco_3750_2 router configuration
- To configure the FortiGate-800 unit
- Cisco_3750_3 router configuration

Cisco_3750_1 router configuration

```
version 12.2
!
hostname Cisco-3750-1
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface Loopback0
   ip address 169.254.100.1 255.255.255.255
!
interface FastEthernet1/0/23
   switchport access vlan 182
   switchport mode access
!
interface FastEthernet1/0/24
   switchport access vlan 172
   switchport mode access
!
interface Vlan172
   ip address 10.31.138.1 255.255.255.0
   ip pim sparse-mode
   ip igmp query-interval 125
   ip mroute-cache distributed
!
interface Vlan182
   ip address 169.254.82.250 255.255.255.0
   ip pim sparse-mode
   ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 169.254.82.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
ip access-list standard Source-RP
   permit 233.254.200.0 0.0.0.255
```

Cisco_3750_2 router configuration

```
version 12.2
!
hostname Cisco-3750-2
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
```

```
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
   switchport access vlan 138
   switchport mode access
!
interface FastEthernet1/0/24
   switchport access vlan 182
   witchport mode access
!
interface Vlan138
   ip address 10.31.138.250 255.255.255.0
   ip pim sparse-mode
   ip mroute-cache distributed
!
interface Vlan182
   ip address 169.254.82.1 255.255.255.0
   ip pim sparse-mode
   ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.138.253
ip route 169.254.100.1 255.255.255.255 169.254.82.250
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255
```

To configure the FortiGate-800 unit

1. Configure the internal and external interfaces.

- **Internal**

  Go to **Network > Interfaces**.

  Select the internal interface.

  Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.138.253 255.255.255.0
Administrative Access:	PING

  Select **OK**.

- **External**

  Go to **Network > Interfaces**.
  Select the external interface.

  Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.130.253 255.255.255.0
Administrative Access:	HTTPS and PING

  Select **OK**.

2. Add a firewall addresses.

   Go to **Policy & Objects > Addresses**.

   - RP

   Select **Create New**.

   Use the following settings:

Category	Address
Name	RP
Type	Subnet
Subnet/IP Range	169.254.100.1/32
Interface	Any
Visibility	<enabled>

   Select **OK**.

   - Multicast source subnet

   Select **Create New**.

   Use the following settings:

Category	Address
Name	multicast_source_subnet

Type	Subnet
Subnet/IP Range	169.254.82.0/24
Interface	Any
Visibility	<enabled>

Select **OK**.

3. Add destination multicast address
   Go to **Policy & Objects > Addresses**.

   Select **Create New**.

   Use the following settings:

Category	Multicast Address
Name	Multicast_stream
Type	Broadcast Subnet
Broadcast Subnet	233.254.200.0/24
Interface	Any
Visibility	<enabled>

   Select **OK**.

4. Add standard security policies to allow traffic to reach the RP.

   Go to **Policy & Objects > IPv4 Policy**.
   - 1st policy

   Select **Create New**.

   Use the following settings:

Incoming Interface	internal
Source Address	all
Outgoing Interface	external
Destination Address	RP
Schedule	always

Service	ALL
Action	ACCEPT

Select **OK**.

- 2nd policy

Select **Create New**

Use the following settings:

Incoming Interface	external
Source Address	RP
Outgoing Interface	internal
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

Select **OK**.

5. Add the multicast security policy.

   Go to **Policy & Objects > Multicast Policy**.

   Select **Create New**.

   Use the following settings:

Incoming Interface	external
Source Address	multicast_source_subnet
Outgoing Interface	internal
Destination Address	multicast_stream
Protocol	Any
Action	ACCEPT

   Select **OK**.

6. Add an access list. (CLI only)

```
config router access-list
```

```
edit Source-RP
   config rule
      edit 1
         set prefix 233.254.200.0 255.255.255.0
         set exact-match disable
      next
   end
```

7. Add some static routes.

   Go to **Network > Static Routes**.

   - Route 1

   Select **Create New**.

   Use the following settings:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	10.31.130.250
Interface	internal
Administrative Distance	<default>
Priority	<default>

   Select **OK**.

   - Route 2

   Select **Create New**.

   Use the following settings:

Destination IP/Mask	169.254.0.0/16
Gateway	10.31.138.250
Interface	external
Administrative Distance	<default>
Priority	<default>

   Select **OK**.

8. Configure multicast routing.

   Go to **Network > Multicast**.

   Add the following Static Rendezvous Point(s):

   - 169.254.100.1

   - Route 1

Select **Create New**.

Use the following settings:

Interface	internal
PIM Mode	Sparse Mode
DR Priority	<not needed in this scenario>
RP Candidate	<not needed in this scenario>
RP Candidate Priority	<not needed in this scenario>

Select **OK.**

- Route 2

Select **Create New**.

Use the following settings:

Interface	external
PIM Mode	Sparse Mode
DR Priority	
RP Candidate	
RP Candidate Priority	

Select **OK**.

Cisco_3750_3 router configuration

```
version 12.2
!
hostname Cisco-3750-3
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
   switchport access vlan 128
   switchport mode access
!
interface FastEthernet1/0/24
```

```
   switchport access vlan 130
   switchport mode access
!
interface Vlan128
   ip address 10.31.128.130 255.255.255.252
   ip pim sparse-mode
   ip mroute-cache distributed
!
interface Vlan130
   ip address 10.31.130.250 255.255.255.0
   ip pim sparse-mode
   ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.130.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255
```

FortiGate PIM-SM debugging examples

Using the example topology shown below, you can trace the multicast streams and states within the three FortiGate units (FGT-1, FGT-2, and FGT-3) using the debug commands described in this section. The command output in this section is taken from the FortiGate unit when the multicast stream is flowing correctly from source to receiver.

PIM-SM debugging topology

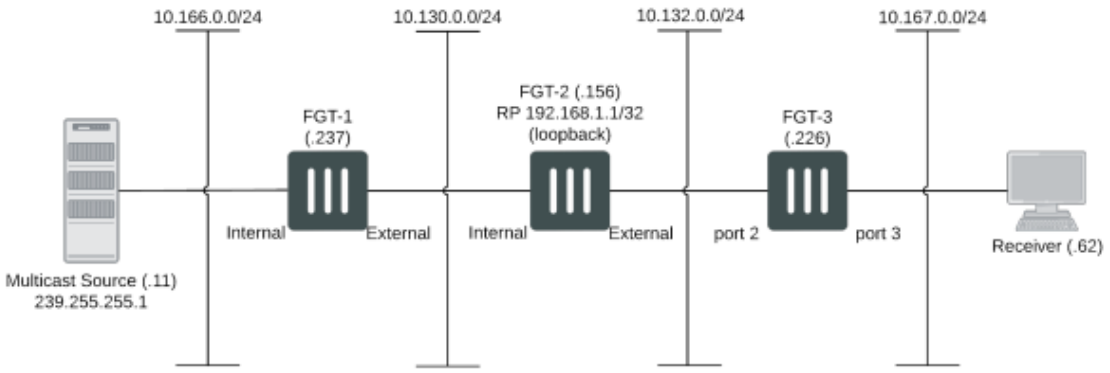

Checking that the receiver has joined the required group

From the last hop router, FGT-3, you can use the following command to check that the receiver has correctly joined the required group.

```
FGT-3 # get router info multicast igmp groups
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
```

```
239.255.255.1 port3 00:31:15 00:04:02 10.167.0.62
```

Only 1 receiver is displayed for a particular group, this is the device that responded to the IGMP query request from the FGT-3. If a receiver is active, the expire time should drop to approximately 2 minutes before being refreshed.

Checking the PIM-SM neighbors

Next, the PIM-SM neighbors should be checked. A PIM router becomes a neighbor when the PIM router receives a PIM hello. Use the following command to display the PIM-SM neighbors of FGT-3:

```
FGT-3 # get router info multicast pim sparse-mode neighbour
Neighbor Interface Uptime/Expires Ver DR
Address Priority/Mode
10.132.0.156 port2 01:57:12/00:01:33 v2 1 /
```

Checking that the PIM router can reach the RP

The rendezvous point (RP) must be reachable for the PIM router (FGT-3) to be able to send the `*,G` join to request the stream. This can be checked for FGT-3 using the following command:

```
FGT-3 # get router info multicast pim sparse-mode rp-mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
RP: 192.168.1.1
Uptime: 07:23:00
```

Viewing the multicast routing table (FGT-3)

The FGT-3 unicast routing table can be used to determine the path taken to reach the RP at 192.168.1.1. You can then check the stream state entries using the following commands:

```
FGT-3 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

(*,*,RP) Entries	This state may be reached by general joins for all groups served by a specified RP.
(*,G) Entries	State that maintains the RP tree for a given group.
(S,G) Entries	State that maintains a source-specific tree for source `S` and group `G`.
(S,G,rpt) Entries	State that maintains source-specific information about source s on the RP tree for `G`. For example, if a source is being received on the source-specific tree, it will normally have been pruned off the RP tree.
FCR	The FCR state entries are for tracking the sources in the <*, G> when <S, G> is not available for any reason, the stream would typically be flowing when this state exists.

Breaking down each entry in detail:

```
(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: JOINED
Local:
port3
Joined:
Asserted:
FCR:
```

The RP will always be listed in a `*,G` entry, the RPF neighbor and interface index will also be shown. In this topology these are the same in all downstream PIM routers. The state is active so the upstream state is joined.

In this case FGT-3 is the last hop router so the IGMP join is received locally on port3. There is no PIM outgoing interface listed for this entry as it is used for the upstream PIM join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.132.0.156
RPF idx: port2
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
Asserted:
Outgoing:
port3
```

This is the entry for the SPT, no RP IS listed. The `S,G` stream will be forwarded out of the stated outgoing interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: NOT PRUNED
Local:
Pruned:
Outgoing:
```

The above `S,G,RPT` state is created for all streams that have both a `S,G` and a `*,G` entry on the router. This is not pruned, in this case, because of the topology: the RP and source are reachable over the same interface.

Although not seen in this scenario, assert states may be seen when multiple PIM routers exist on the same LAN, which can lead to more than one upstream router having a valid forwarding state. Assert messages are used to elect a single forwarder from the upstream devices.

Viewing the PIM next-hop table

The PIM next-hop table is also very useful for checking the various states, it can be used to quickly identify the states of multiple multicast streams.

```
FGT-3 # get router info multicast pim sparse-mode next-hop
Flags: N = New, R = RP, S = Source, U = Unreachable
Destination Type Nexthop Nexthop Nexthop Metric Pref Refcnt
Num Addr Ifindex
_____________________________________________________________
10.166.0.11 ..S. 1 10.132.0.156 9 21 110 3
192.168.1.1 .R.. 1 10.132.0.156 9 111 110 2
```

Viewing the PIM multicast forwarding table

Also, you can check the multicast forwarding table showing the ingress and egress ports of the multicast stream.

```
FGT-3 # get router info multicast table

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL threshold)

(10.166.0.11, 239.255.255.1), uptime 04:02:55, stat expires 00:02:25
Owner PIM-SM, Flags: TF
Incoming interface: port2
Outgoing interface list:
port3 (TTL threshold 1)
```

Viewing the kernel forwarding table

Also, the kernel forwarding table can be verified, however this should give similar information to the above command:

```
FGT-3 # diag ip multicast mroute
grp=239.255.255.1 src=10.166.0.11 intf=9 flags=(0x10000000)[ ] status=resolved
last_assert=2615136 bytes=1192116 pkt=14538 wrong_if=0 num_ifs=1
index(ttl)=[6(1),]
```

Viewing the multicast routing table (FGT-2)

If you check the output on FGT-2, there are some small differences:

```
FGT-2 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
Local:
Joined:
external
Asserted:
FCR:
```

The `*,G` entry now has a joined interface rather than local because it has received a PIM join from FGT-3 rather than a local IGMP join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.130.0.237
RPF idx: internal
SPT bit: 1
Upstream State: JOINED
```

```
Local:
Joined:
external
Asserted:
Outgoing:
external
```

The `S,G` entry shows that we have received a join on the external interface and the stream is being forwarded out of this interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: PRUNED
Local:
Pruned:
Outgoing:
External
```

The `S,G,RPT` is different from FGT-3 because FGT-2 is the RP, it has pruned back the SPT for the RP to the first hop router.

Viewing the multicast routing table (FGT-1)

FGT-1 again has some differences with regard to the PIM-SM states. There is no `*,G` entry because it is not in the path of a receiver and the RP.

```
FGT-1_master # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

Below the `S,G` is the SPT termination because this FortiGate unit is the first hop router. The RPF neighbor always shows as 0.0.0.0 because the source is local to this device. Both the joined and outgoing fields show as external because the PIM join and the stream is egressing on this interface.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 0.0.0.0
RPF idx: None
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
external
Asserted:
Outgoing:
external
```

The stream has been pruned back from the RP because the end-to-end SPT is flowing. In this case, there is no requirement for the stream to be sent to the RP.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 0.0.0.0
RPF nbr: 10.130.0.156
RPF idx: external
Upstream State: RPT NOT JOINED
```

```
Local:
Pruned:
Outgoing:
```

Example multicast DNAT configuration

The example topology shown and described below shows how to configure destination NAT (DNAT) for two multicast streams. Both of these streams originate from the same source IP address, which is 10.166.0.11. The example configuration keeps the streams separate by creating 2 multicast NAT policies.

In this example, the FortiGate units have the following roles:

- FGT-1 is the RP for dirty networks, 233.0.0.0/8.
- FGT-2 performs all firewall and DNAT translations.
- FGT-3 is the RP for the clean networks, 239.254.0.0/16.
- FGT-1 and FGT-3 are functioning as PM enabled routers and could be replaced can be any PIM enabled router.

This example only describes the configuration of FGT-2.

FGT-2 performs NAT so that the receivers connected to FGT-3 receive the following translated multicast streams.

- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.2.2.1, FGT-3 translates the source and destination IPs to 192.168.20.1 and 239.254.1.1
- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.3.3.1, FGT-3 translates the source and destination IPs to 192.168.20.10 and 239.254.3.1

Example multicast DNAT topology

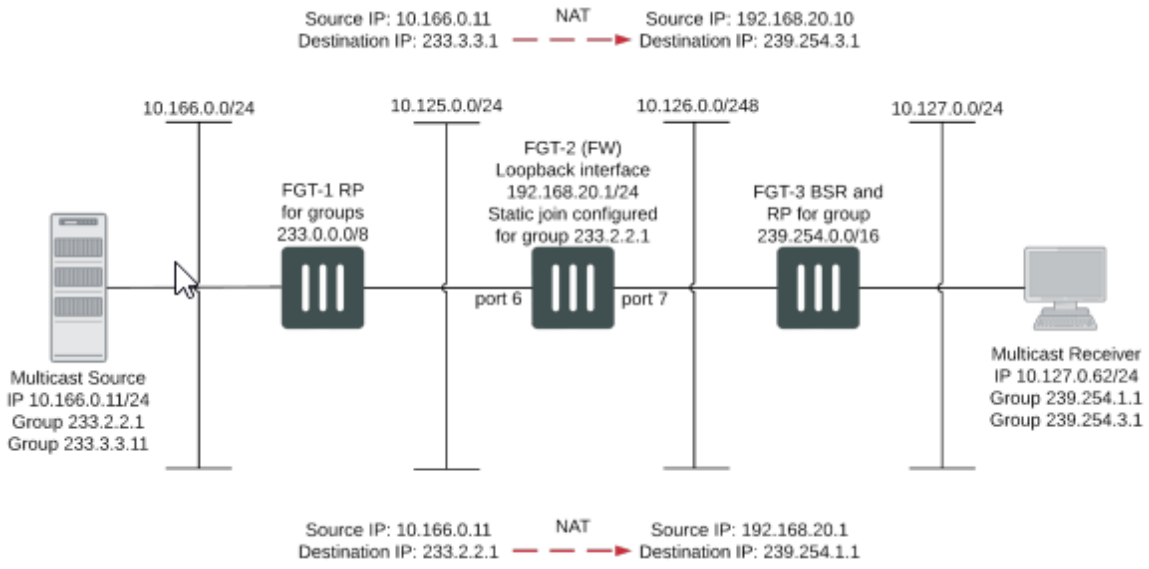

To configure FGT-2 for DNAT multicast

1. Add a loopback interface. In the example, the loopback interface is named `loopback`.

```
config system interface
```

```
      edit loopback
         set vdom root
         set ip 192.168.20.1 255.255.255.0
         set type loopback
      next
   end
```

2. Add PIM and add a unicast routing protocol to the loopback interface as if it was a normal routed interface. Also, add static joins to the loopback interface for any groups to be translated.

```
   config router multicast
      config interface
         edit loopback
            set pim-mode sparse-mode
               config join-group
                  edit 233.2.2.1
                  next
                  edit 233.3.3.1
                  next
               end
         next
```

3. In this example, to add firewall multicast policies, different source IP addresses are required so you must first add an IP pool:

```
   config firewall ippool
      edit Multicast_source
         set endip 192.168.20.20
         set interface port6
         set startip 192.168.20.10
      next
   end
```

4. Add the translation security policies.

   Policy 2, which is the source NAT policy, uses the actual IP address of port6. Policy 1, the DNAT policy, uses an address from the IP pool. The source and destination addresses will need to be previously created address objects. For this example, 233.3.3.1 255.255.255.255 will be represented by "example-addr_1" and 10.166.0.11 255.255.255.255 will be represented by "example-addr_2". You will likely want to use something more intuitive from your own network.

```
   config firewall multicast-policy
      edit 1
         set dnat 239.254.3.1
         set dstaddr example-addr_1
         set dstintf loopback
         set nat 192.168.20.10
         set srcaddr example-addr_2
         set srcintf port6
      next
      edit 2
         set dnat 239.254.1.1
         set dstaddr 233.2.2.1 255.255.255.255
         set dstintf loopback
         set nat 192.168.20.1
         set srcaddr 10.166.0.11 255.255.255.255
         set srcintf port6
      next
   end
```

5. Add a firewall multicast policy to forward the stream from the loopback interface to the physical outbound interface.

   This example is an any/any policy that makes sure traffic accepted by the other multicast policies can exit the FortiGate unit.

```
config firewall multicast-policy
   edit 3
      set dstintf port7
      set srcintf loopback
   next
end
```

Example PIM configuration that uses BSR to find the RP

This example shows how to configure a multicast routing network for a network consisting of four FortiGate-500A units (FortiGate-500A_1 to FortiGate-550A_4). A multicast sender is connected to FortiGate-500A_2. FortiGate-500A_2 forwards multicast packets in two directions to reach Receiver 1 and Receiver 2.

The configuration uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface `lo0` must join the 236.1.1.1 group (source).

This example describes:

- Commands used in this example
- Configuration steps
- Example debug commands

PIM network topology using BSR to find the RP

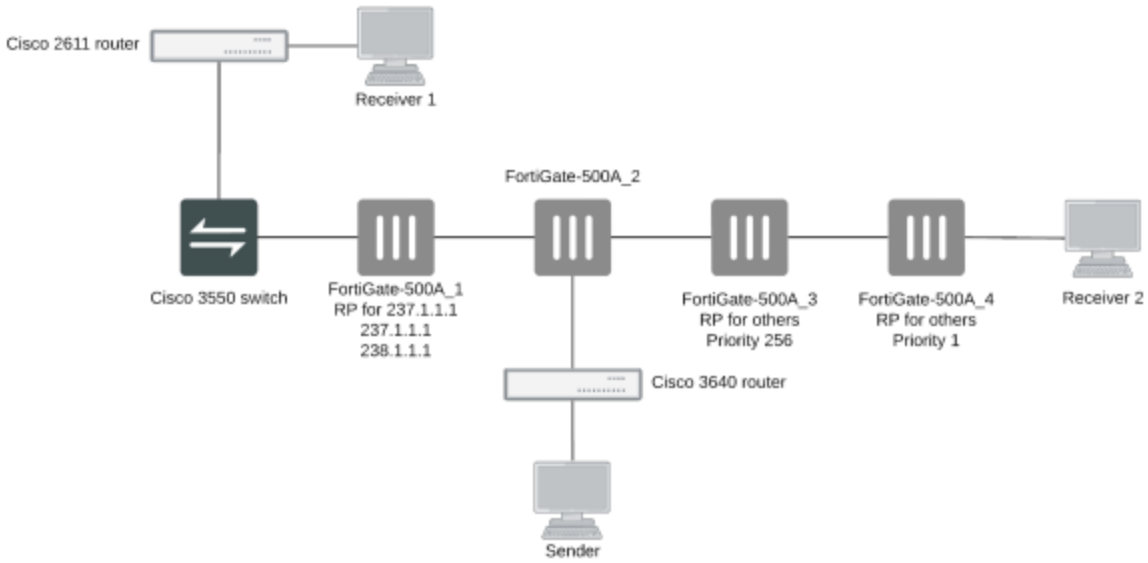

Commands used in this example

This example uses CLI commands for the following configuration settings:

- Adding a loopback interface (lo0)
- Defining the multicast routing
- Adding the NAT multicast policy

Adding a loopback interface

Where required, the following command is used to define a loopback interface named `lo0`.

```
config system interface
   edit lo0
      set vdom root
      set ip 1.4.50.4 255.255.255.255
      set allowaccess ping https ssh snmp http telnet
      set type loopback
   next
end
```

Defining the multicast routing

In this example, the following command syntax is used to define multicast routing.

The example uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface `lo0` must join the 236.1.1.1 group (source).

```
config router multicast
   config interface
      edit port6
         set pim-mode sparse-mode
      next
      edit port1
         set pim-mode sparse-mode
      next
      edit lo0
         set pim-mode sparse-mode
         set rp-candidate enable
            config join-group
               edit 236.1.1.1
               next
            end
         set rp-candidate-priority 1
      next
   end
set multicast-routing enable
   config pim-sm-global
      set bsr-allow-quick-refresh enable
      set bsr-candidate enable
      set bsr-interface lo0
      set bsr-priority 200
   end
end
```

Adding the NAT multicast policy

In this example, the incoming multicast policy does the address translation.

The NAT address should be the same as the IP address of the of loopback interface. The DNAT address is the translated address, which should be a new group.

```
config firewall multicast-policy
   edit 1
      set dstintf port6
      set srcintf lo0
   next
   edit 2
      set dnat 238.1.1.1
      set dstintf lo0
      set nat 1.4.50.4
      set srcintf port1
   next
```

Configuration steps

In this sample, FortiGate-500A_1 is the RP for the group 228.1.1.1, 237.1.1.1, 238.1.1.1, and FortiGate-500A_4 is the RP for the other group which has a priority of1. OSPF is used in this example to distribute routes including the loopback interface. All firewalls have full mesh security policies to allow any to any.

- In the FortiGate-500A_1 configuration, the NAT policy translates source address 236.1.1.1 to 237.1.1.1
- In the FortiGate-500A_4 configuration, the NAT policy translates source 236.1.1.1 to 238.1.1.1
- Source 236.1.1.1 is injected into network as well

The following procedures include the CLI commands for configuring each of the FortiGate units in the example configuration.

To configure FortiGate-500A_1

1. Configure multicast routing:

```
config router multicast
   config interface
      edit port5
         set pim-mode sparse-mode
      next
      edit port4
         set pim-mode sparse-mode
      next
      edit lan
         set pim-mode sparse-mode
      next
      edit port1
         set pim-mode sparse-mode
      next
      edit lo999
         set pim-mode sparse-mode
      next
      edit lo0
         set pim-mode sparse-mode
         set rp-candidate enable
         set rp-candidate-group 1
      next
   end
set multicast-routing enable
   config pim-sm-global
      set bsr-candidate enable
      set bsr-interface lo0
   end
```

```
end
```

2. Add multicast security policies:

```
config firewall multicast-policy
   edit 1
      set dstintf port5
      set srcintf port4
   next
   edit 2
      set dstintf port4
      set srcintf port5
   next
   edit 3
   next
end
```

3. Add router access lists:

```
config router access-list
   edit 1
      config rule
         edit 1
            set prefix 228.1.1.1 255.255.255.255
            set exact-match enable
         next
         edit 2
            set prefix 237.1.1.1 255.255.255.255
            set exact-match enable
         next
         edit 3
            set prefix 238.1.1.1 255.255.255.255
            set exact-match enable
         next
      end
   next
end
```

To configure FortiGate-500A_2

1. Configure multicast routing:

```
        config router multicast
           config interface
              edit "lan"
                 set pim-mode sparse-mode
next
edit "port5"
  set pim-mode sparse-mode
              next
              edit "port2"
                 set pim-mode sparse-mode
              next
              edit "port4"
                 set pim-mode sparse-mode
              next
              edit "lo_5"
                 set pim-mode sparse-mode
                    config join-group
```

```
                edit 236.1.1.1
                next
            end
        next
    end
    set multicast-routing enable
end
```

2. Add multicast security policies:

```
config firewall multicast-policy
    edit 1
        set dstintf lan
        set srcintf port5
    next
    edit 2
        set dstintf port5
        set srcintf lan
    next
    edit 4
        set dstintf lan
        set srcintf port2
    next
    edit 5
        set dstintf port2
        set srcintf lan
    next
    edit 7
        set dstintf port1
        set srcintf port2
    next
    edit 8
        set dstintf port2
        set srcintf port1
    next
    edit 9
        set dstintf port5
        set srcintf port2
    next
    edit 10
        set dstintf port2
        set srcintf port5
    next
    edit 11
        set dnat 237.1.1.1
        set dstintf lo_5
        set nat 5.5.5.5
        set srcintf port2
    next
    edit 12
        set dstintf lan
        set srcintf lo_5
    next
    edit 13
        set dstintf port1
        set srcintf lo_5
    next
    edit 14
```

```
      set dstintf port5
      set srcintf lo_5
   next
   edit 15
      set dstintf port2
      set srcintf lo_5
   next
   edit 16
   next
end
```

To configure FortiGate-500A_3

1. Configure multicast routing:

```
config router multicast
   config interface
      edit port5
         set pim-mode sparse-mode
      next
      edit port6
         set pim-mode sparse-mode
      next
      edit lo0
         set pim-mode sparse-mode
         set rp-candidate enable
         set rp-candidate-priority 255
      next
      edit lan
         set pim-mode sparse-mode
      next
   end
set multicast-routing enable
   config pim-sm-global
      set bsr-candidate enable
      set bsr-interface lo0
   end
end
```

2. Add multicast security policies:

```
config firewall multicast-policy
   edit 1
      set dstintf port5
      set srcintf port6
   next
   edit 2
      set dstintf port6
      set srcintf port5
   next
   edit 3
      set dstintf port6
      set srcintf lan
   next
   edit 4
      set dstintf lan
      set srcintf port6
   next
   edit 5
```

```
            set dstintf port5
            set srcintf lan
         next
         edit 6
            set dstintf lan
            set srcintf port5
         next
      end
```

To configure FortiGate-500A_4

1. Configure multicast routing:

```
      config router multicast
         config interface
            edit port6
               set pim-mode sparse-mode
            next
            edit lan
               set pim-mode sparse-mode
            next
            edit port1
               set pim-mode sparse-mode
            next
            edit lo0
               set pim-mode sparse-mode
               set rp-candidate enable
                  config join-group
                     edit 236.1.1.1
                     next
                  end
               set rp-candidate-priority 1
            next
         end
      set multicast-routing enable
         config pim-sm-global
set bsr-allow-quick-refresh enable
            set bsr-candidate enable
            set bsr-interface lo0
            set bsr-priority 1
         end
      end
```

2. Add multicast security policies:

```
      config firewall policy
         edit 1
            set srcintf lan
            set dstintf port6
            set srcaddr all
            set dstaddr all
            set action accept
            set schedule always
            set service ANY
         next
         edit 2
            set srcintf port6
            set dstintf lan
            set srcaddr all
```

```
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 3
      set srcintf port1
      set dstintf port6
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 4
      set srcintf port6
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 5
      set srcintf port1
      set dstintf lan
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 6
      set srcintf lan
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 7
      set srcintf port1
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
   next
   edit 8
      set srcintf port6
      set dstintf lo0
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
```

```
        set service ANY
      next
      edit 9
        set srcintf port1
        set dstintf lo0
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
      next
      edit 10
        set srcintf lan
        set dstintf lo0
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
      next
    end
```

Troubleshooting

Netflow support

Netflow is a networking feature, introduced by Cisco, to collect and export information about traffic flow through routers. IPFIX (Internet Protocol Flow Information Export) is the standardized Internet Protocol based on NetFlow version 9. The standards requirements for IPFIX are outlined in RFC 3197, and its basic specifications and other information are documented in RFC 5103, RFC 6759, and RFC 7011 through RFC 7015.

You can enable and configure NetFlow traffic, using the following CLI commands:

```
config system netflow
   set collector-ip <collector IP>
   set collector-port <NetFlow collector port>
   set csource-ip <Source IP for NetFlow agent>
   set cactive-flow-timeout <time in minutes of timeout to report active flows>
   set cinactive-flow-timeout <time in seconds of timeout for periodic report of finished
       flows>
   end
```

You can also configure these settings per VDOM, using the following CLI command:

```
config system vdom-netflow
```

You must also enable a Netflow sampler on specific interfaces.

sFlow support

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. FortiOS implements sFlow version 5.

sFlow uses packet sampling to monitor network traffic. The sFlow Agent captures packet information at defined intervals and sends them to an sFlow Collector for analysis, providing real-time data analysis. The information sent is only a sampling of the data for minimal impact on network throughput and performance.

The sFlow Agent is embedded in the FortiGate unit. Once configured, the FortiGate unit sends sFlow datagrams of the sampled traffic to the sFlow Collector, also called an sFlow Analyzer. The sFlow Collector receives the datagrams, and provides real-time analysis and graphing to indicate where potential traffic issues are occurring. sFlow Collector software is available from a number of third party software vendors.

sFlow data captures only a sampling of network traffic, not all traffic like the traffic logs on the FortiGate unit. Sampling works by the sFlow Agent looking at traffic packets when they arrive on an interface. A decision is made whether the packet is dropped and allowed to be to its destination or if a copy is forwarded to the sFlow Collector. The sample used and its frequency are determined during configuration.

sFlow is not supported on virtual interfaces such as vdom link, ipsec, ssl.root or gre.

The sFlow datagram sent to the Collector contains the information:

- Packet header (e.g. MAC,IPv4,IPv6,IPX,AppleTalk,TCP,UDP, ICMP)
- Sample process parameters (rate, pool etc.)
- Input/output ports
- Priority (802.1p and TOS)
- VLAN (802.1Q)
- Source/destination prefix
- Next hop address
- Source AS, Source Peer AS
- Destination AS Path
- Communities, local preference
- User IDs (TACACS/RADIUS) for source/destination
- URL associated with source/destination
- Interface statistics (RFC 1573, RFC 2233, and RFC 2358)

sFlow agents can be added to any type of FortiGate interface. sFlow isn't supported on some virtual interfaces such as VDOM link, IPsec, gre, and ssl.root.

For more information on sFlow, Collector software and sFlow MIBs, visit www.sflow.org.

Configuration

sFlow configuration is available only from the CLI. Configuration requires two steps: enabling the sFlow Agent and configuring the interface for the sampling information.

Enable sFlow

```
config system sflow
   set collector-ip <ip_address>
   set collector-port <port_number>
   set source-ip <ip_address>
end
```

The default port for sFlow is UDP 6343. To configure in VDOM, use the following CLI commands:

```
config system vdom-sflow
   set vdom-sflow enable
   set collector-ip <ip_address>
   set collector-port <port_number>
   set source-ip <ip_address>
end
```

Configure sFlow agents per interface.

```
config system interface
   edit <interface_name>
      set sflow-sampler enable
      set sample-rate <every_n_packets>
      set sample-direction [tx | rx | both]
      set polling-interval <seconds>
   end
```

Copyright© (Undefined variable: FortinetVariables.CopyrightYear) Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.